

C9

```
Childbone 138 (RSLauxSupp.c)
DebugLogFds 122 (RSLwsvr.c)
DemuxAuxChildren 120 (RSLauxmain.c)
DetermineGlobalDriveUse 119 (RSLwsvr.c)
ForwardXcpioGenProgress 118 (RSLauxmain.c)
GetAuxprocResults 53 (RSLauxmgr.c)
HandleWorkItemRestoreResults 13 (RSLwsvr.c)
InitiateWorkItemRestore 11 (RSLwsvr.c)
InterpretWorkItemRestoreResults 21 (RSLwsvr.c)
KillWorkItemRestore 55 (RSLauxmgr.c)
QuitWorkItemRestore 51 (RSLauxmgr.c)
RunWorkItemRestores 3 (RSLwsvr.c)
RunWorkItemRestoresForTrail 17 (RSLwsvr.c)
Select 10 (RSLwsvr.c)
SendRunningWorkItemsQuit 20 (RSLwsvr.c)
StartWorkItemRestore 46 (RSLauxmgr.c)
StartupAuxprocess 31 (RSLauxmgr.c)
auxcmdpacket 133 (RSLauxSupp.c)
auxproc_comm_warning 136 (RSLauxSupp.c)
auxres2 134 (RSLauxSupp.c)
auxresults 135 (RSLauxSupp.c)
decodecookie 96 (RSLauxmain.c)
do_auxproc 62 (RSLauxmain.c)
ebr_direct_rcmd 105 (RSLauxmain.c)
fd_avail_1_wait_intr 102 (RSLauxmain.c)
fd_avail_test 137 (RSLauxSupp.c)
fd_avail_test1 104 (RSLauxmain.c)
fd_avail_test_intr 103 (RSLauxmain.c)
generate_rcmdpath 44 (RSLauxmgr.c)
looprw 125 (RSLauxSupp.c)
main 61 (RSLauxmain.c)
make_remote_cpioGen_cmd 42 (RSLauxmgr.c)
parse_remote_stderr_info2 112 (RSLauxmain.c)
pread_or_die 129 (RSLauxSupp.c)
pread_or_warn 130 (RSLauxSupp.c)
pwrite_or_die 131 (RSLauxmain.c)
pwrite_or_warn 132 (RSLauxSupp.c)
rb_getmethod 110 (RSLauxmain.c)
read_CD_L_no_eintr 100 (RSLauxmain.c)
read_no_eintr 127 (RSLauxSupp.c)
recover_size_prefix 68 (RSLauxmain.c)
reset_recovery_privileges 50 (RSLauxmgr.c)
set_recovery_privileges 48 (RSLauxmgr.c)
sigterm_handler 98 (RSLauxmain.c)
sigusr1_handler 97 (RSLauxmain.c)
start_cpioGen 34 (RSLauxmgr.c)
test_fd 23 (RSLwsvr.c)
test_fd_hup 24 (RSLwsvr.c)
write_CD_L_no_eintr 101 (RSLauxmain.c)
write_no_eintr 128 (RSLauxSupp.c)
z_exec_separate_auxproc 99 (RSLauxmain.c)
z_rcmdfilter 69 (RSLauxmain.c)
```


RSIwtsvr.c

1

```
DebugLogFds.....22
DetermineGlobalDriveUse 19
HandleWorkItemRestoreResults...13
InitiateWorkItemRestore 11
InterpretWorkItemRestoreResults...21
RunWorkItemRestores 3
RunWorkItemRestoresForTrail...17
Select 10
SendRunningWorkItemsQuit...20
test_fd 23
test_fd_hup.....24
```

RSIauxmgr.c

29

```
GetAuxprocResults.....53
KillWorkItemRestore 55
QuitWorkItemRestore.....51
StartWorkItemRestore 46
StartupAuxprocess.....31
generate_rcmdpath 44
make_remote_cpiogen_cmd...42
reset_recovery_privileges 50
set_recovery_privileges...48
start_cpiogen 34
```

RSIauxmain.c

57

```
DemuxAuxChildren 120
ForwardXcpiogenProgress...118
decodecookie 96
do_auxproc.....62
ebr_direct_rcmd 105
fd_avail_1_wait_intr....102
fd_avail_test1 104
fd_avail_test_intr.....103
main 61
parse_remote_stderr_info2..112
rb_getmethod 110
read_CDL_no_eintr.....100
recover_size_prefix 68
sigterm_handler.....98
sigusr1_handler 97
write_CDL_no_eintr.....101
z_exec_separate_auxproc 99
z_rcmdfilter.....69
```

RSIauxsupp.c

125

```
ChildDone.....138
auxcmdpacket 133
auxproc_comm_warning.....136
auxres2 134
auxresults.....135
fd_avail_test 137
looprw.....125
pread_or_warn 129
pwrite_or_die 130
pwrite_or_warn.....131
read_no_eintr 127
write_no_eintr.....128
```



```

1  /*****
2  **
3  ** File Name:   RSLwsvr.c
4  **
5  ** Copyright (c) 1998,1999 by EMC Corporation.
6  **
7  ** Purpose:
8  ** -----
9  ** The intent of the contents of this file is to implement the
10 ** functions the control execution of work item restores for
11 ** Restore
12 ** Service Library.
13 **
14 ** These functions are provided to allow:
15 **
16 ** The following functions comprise restoral management:
17 **
18 ** RunWorkItemRestores()
19 **
20 **
21 ** Compile-Time Options:
22 ** This section must list any compile time definitions
23 ** which will affect this header.
24 **
25 *****/
26
27 #define _POSIX_SOURCE 1
28
29 /*
30 * System headers.
31 */
32
33 #include <sys/time.h>
34 #include <sys/types.h>
35 #include <sys/wait.h>
36 #include <values.h>
37
38 /*
39 * Epoch headers.
40 */
41
42 #include <eb/eb_port.h>
43 #include <eb/rb_log.h>
44 #include <ebutil/ebutil.h>
45 #include <restore/RDprogmsg.h>
46
47 /*
48 * Local headers
49 */
50
51 #include <RSLspxlts.h>
52 #include <RSLremfd.h>
53 #include <EDMRESchedApi.h>
54 #include <EDMREHandleMgrApi.h>
55 #include <RSLIntern.h>
56 #include <RSLrbrmain.h>
57 #include <restore/EDMRESubmitApi.h>
58 #include <EDMREDrainApi.h>
59
60 #define STR_SURE(str) (str) ? str:""
61
62 /*
63 */

```

```

64 * RunWorkItemRestores
65 * {
66 *
67 * Set the number of drives being used for the life of the restore.
68 * SetQuitFlag = FALSE
69 * TrailRestoresLeft = { # of trail restores }
70 * TrailRestoresRunning = 0;
71 *
72 * while drive available.
73 * RunTrail -- Set drive concurrency for trail restore.
74 * TrailRestoresRunning++;
75 * while OKToRunWIForTrail
76 * StartWIRestore
77 *
78 * while(1)
79 * if(QuitTest)
80 * Send WICancelis
81 * SetQuitFlag = TRUE
82 * Select(from_pipe, timeout (5 seconds))
83 * for each WI that completes
84 * interperate return.
85 * Drain progress.
86 * Send final progress for work item.
87 * if(WIFailed)
88 * if(OKToReschedule && SetQuitFlag == FALSE)
89 * Add to TrailQueue
90 * if(TrailRestoresAsMoreWorkItems && SetQuitFlag == FALSE)
91 * while OKToRunWIForTrail
92 * StartWIRestore
93 * else -- RunNewTrailRestore
94 * EndTrailRestore(prevTrailQueue)
95 * TrailRestoresRunning--;
96 * TrailRestoresLeft--;
97 * while (drive available &&
98 * SetQuitFlag == FALSE &&
99 * TrailRestoresLeft)
100 *
101 * RunTrail -- Set drive concurrency for trail restore.
102 * TrailRestoresRunning++;
103 * while (
104 * OKToRunWIForTrail && SetQuitFlag == FALSE)
105 * StartWIRestore
106 * end while
107 * end while
108 *
109 * End for each WI completes
110 * if (((SetQuitFlag == TRUE) && (TrailRestoresRunning == 0)) ||
111 * TrailRestoresLeft == 0)
112 * return; exit loop.
113 * end while(1)
114 * }
115 *
116 */
117
118 /* Scrubs */
119 static int
120 InterpretWorkItemRestoreResults(wi_restore_results *results);
121
122 static int
123 test_fd(int fd);
124
125 static void
126 DebugLogFds(char *error_msg,
127 fd_set *fds);

```

```

129 static int
130 DetermineGlobalDriveUse();
131
132 static int
133 test_fd_hup(int fd);
134
135 static int
136 FindTrailIDForWorkItem(int handle,
137 int *trailID,
138 int *status);
139
140 static int
141 SendRunningWorkItemsQuit();
142
143 /* End Stubs */
144
145 static int
146 Select(int nfds,
147 fd_set *readfds,
148 fd_set *writefds,
149 fd_set *exceptfds,
150 struct timeval *timeout);

```

```

153 static int
154 HandleWorkItemRestoreResults(int FromFD,
155 int *trailID,
156 wi_restore_results *results);
157
158 static int
159 RunWorkItemRestoresForTrail(const int TrailID,
160 const int CountDrivesAvailable,
161 boolean_ty (*CancelRestoreTest)(),
162 boolean_ty *QuitFlag,
163 int *CountDrivesInUse);

```

```

164 /*
165 * RunWorkItemRestores()
166 *
167 * Runs a set of work item restores.
168 *
169 * Args:
170 * SubmitObject
171 * CancelRestoreTest().
172 *
173 * Returns: int 0 for success.
174 */
175 int
176 RunWorkItemRestores(int SubmitObject, boolean_ty (*CancelRestoreTest)()

```

```

177 {
178     boolean_ty QuitFlag = FALSE; /* Has the user requested a quit. */
179     boolean_ty SentQuit = FALSE; /* Have we initiated the quit. */

```

```

181     int TrailRestoresRunning = 0;
182     /* The number of trail restores running. */
183     int TrailRestoresLeft;
184     /* The number of trail restores left. */
185     int TrailRestoresTotal;
186     /* The number of trail restores total. */

```

```

188     int CountDrivesAvailable; /* The count of drives available. */
189     int CountDrivesInUse = 0; /* The count of drives in use. */

```

```

191     int temp_status;
192     int HighestActiveTrail = 0;
193     /* The trail queues are ordered from 1 to n. */

```

```

190     /* This is the highest trail running */

```

```

192     if(debugmode)
193     {
194         (void)rbe_user_error(0,
195         "DEBUG: Running RunWorkItemRestores.");

```

```

196     }
197     /* GenerateTrailQueues()
198     * Buckets the work items into trail queues.
199     * The trail queues are sorted
200     * in the order which the restores should run.

```

```

201     */
202     if(0 != GenerateTrailQueues(SubmitObject,
203     &TrailRestoresTotal,
204     &temp_status))

```

```

205     {
206         (void)rbe_user_error(0,
207         "Internal error: Cannot generate trail
208         queues, cannot continue.");

```

```

209     return -1;
210 }
211 TrailRestoresLeft = TrailRestoresTotal;

```

```

212 CountDrivesAvailable = DetermineGlobalDriveUse(&SubmitObject*);
213
214 if(debugmode)

```

```

215 {
216     (void)rbe_user_error(0,
217     "DEBUG: RunWorkItemRestores for %d trails.",
218     TrailRestoresTotal);

```

```

219 }
220 /*
221 * This is the start up loop to get the initial work item
222 * restores started.
223 */

```

```

224 QuitFlag = CancelRestoreTest();
225 while((CountDrivesInUse < CountDrivesAvailable) &&
226 (HighestActiveTrail < TrailRestoresTotal) &&
227 (FALSE == QuitFlag))

```

```

228 {
229     int submitObjID = 0;
230     int submitElemID = 0;

```

```

231     HighestActiveTrail++;
232     /* Activate the Trail Queue.
233     * This allows the trail queues to be used to
234     * determine the work item restores to run.

```

```

235     if(0 != ActivateTrailQueue(HighestActiveTrail,
236     1,
237     &temp_status))

```

```

238     {
239         (void)rbe_user_error(0,
240         "Internal error: Cannot activate trail
241         queues(1) for trailid %d, cannot continue.",
242         HighestActiveTrail);

```

```

243     return -1;
244 }
245

```

```

251 2 /*
252 2 * This sets the number of drives and media access concurrency for
253 2 * i.e. The count of running work item restores for this trail.
254 2 * Today this is one.
255 2 */
256 3 if(0 != SetIODrivesAcquired(HighestActiveTrail, 1, &temp_status))
257 3 {
258 3     (void)rbe_user_error(0,
259 3         "Internal error: Cannot set drive acquired(
260 2         1) for trailid %d, cannot continue.", HighestActiveTrail);
261 2     return -1;
262 2 }
263 2 if (0 > (temp_status = RunWorkItemRestoresForTrail(
264 2     HighestActiveTrail,
265 2     CountDrivesAvailable,
266 2     CancelRestoreTest,
267 2     &QuitFlag,
268 2     &CountDrivesInUse)))
269 3 {
270 3     /* RunWorkItemRestoresForTrail does its own error logging. */
271 2     return -1;
272 2 }
273 2 if(temp_status == 0)
274 3 {
275 3     (void)rbe_log_stats(0,
276 3         "Trail %d restore had no work item to run(
277 3         1).", HighestActiveTrail);
278 2     /* more work may be need to recover from this error condition. */
279 2 }
280 3 if(temp_status > 0)
281 3 {
282 2     TrailRestoresRunning++;
283 1 } /* End while() initial startup loop */
284 1 while(1)
285 2 {
286 2     int HighestFd = 0;
287 2     fd_set WorkItemFromFds;
288 2     int retStatus;
289 2     struct timeval timeout = (5, 0);
290 2
291 2     if((QuitFlag) && (!SentQuit))
292 2     {
293 3         (void)rbe_log_stats(0,
294 3             "Restore was quit by user. Quitting restore,
295 3             this could take a while.");
296 2     }
297 3     SendRunningWorkItemsQuit();
298 3     SentQuit = TRUE;
299 2 }
300 2 if(0 != getFromSet(&WorkItemFromFds, &HighestFd, &retStatus))
301 3 {
302 3     (void)rbe_user_error(0,
303 3         "Internal error: Cannot get auxproc result
304 3         fds, cannot continue.");

```

```

306 3     return -1;
307 2 }
308 2 #if 0
309 2     if(debugmode)
310 2     {
311 3         DebugLogFds("The file descriptors to wait on are ",
312 3             &WorkItemFromFds);
313 2     }
314 2 #endif
315 2 if(0 > (retStatus = Select(HighestFd + 1,
316 2     &WorkItemFromFds,
317 2     NULL, NULL,
318 2     &timeout)))
319 3 {
320 3     /* error */
321 3     (void)rbe_user_error(RBRECOVER_MKERR(errno),
322 3         "Internal error: Cannot get auxproc result
323 3         fds, cannot continue.");
324 2 }
325 2 return -1;
326 2 }
327 2 else if (0 == retStatus)
328 3 { /* timed out */
329 3     QuitFlag = CancelRestoreTest();
330 2 }
331 2 else
332 3 { /* Available fds */
333 2     int ReadyFds = retStatus;
334 2     int FoundFds = 0;
335 2     int index;
336 3     if(debugmode)
337 3     {
338 3         DebugLogFds("The file descriptors ready to read are ",
339 3             &WorkItemFromFds);
340 2     }
341 2     /*
342 2     * If there are available fds then we may want to
343 2     * schedule the next work item restore. We should
344 2     * check if the user initiated a quit.
345 2     */
346 2     QuitFlag = CancelRestoreTest();
347 2     for(index = 0;
348 2         (index < (HighestFd + 1)) && (FoundFds < ReadyFds);
349 2         index++)
350 2     {
351 3         int StartWorkItemForTrail = 0;
352 3         if(FD_ISSET(index, &WorkItemFromFds))
353 3         {
354 4             int TrailID;
355 4             int TrailAcquired;
356 4             w1_restore_results results;
357 4             FoundFds++;
358 4             memset(&results, 0, sizeof(w1_restore_results));
359 4             if(0 != HandleWorkItemRestoreResults(index,
360 4                 &TrailID,
361 4                 &results))
362 4             {
363 5
364 5
365 5
366 5
367 5
368 5
369 5
370 6

```



```

371 6      /* HandleWorkItemRestoreResults will do its own logging! */
372 6      return -1;
373 5  }
374 5  /*
375 5  * This is where we may want to retry the Work Item
376 5  * Based on if it passes or fails
377 5  */
378 5  CountDrivesInUse--;

379 5

382 5  if (0 > (StartWorkItemForTrail =
383 5      RunWorkItemRestoresForTrail(TrailID,
384 5      CountDrivesAvailable,
385 5      CancelRestoreTest,
386 5      &QuitFlag,
387 5      &CountDrivesInUse)))
388 6  {
389 6      /* RunWorkItemRestoresForTrail does its own logging. */
390 6      return -1;
391 5  }

393 5  else if (StartWorkItemForTrail == 0)
394 5  {
395 5      /* 0 work items started above,
396 5      * Lets check to see if this is the last work item
397 5      * this trail
398 5      */
399 6      int wiCount;

401 6      if (0 != GetRunningWI(TrailID, &wiCount, &temp_status))
402 7      {
403 7          (void)rbe_user_error(0,
404 7              "Internal error: Cannot determine
              number of running work items for trail, cannot continue.");
405 7      }

406 7      return -1;
407 6  }
408 6  if (debugmode)
409 7  {
410 7      (void)rbe_user_error(0,
411 7          "DEBUG: RunWorkItemRestores no
              more work items left for trailid '%d',
              but %d wiCount workitem still running.",
              TrailID, wiCount);
412 7  }

413 7  /* Testing for No work items left running or started
414 6  * For this trail.
415 6  */
416 6  if ((0 == wiCount) && (0 == StartWorkItemForTrail))
417 7  {
418 6      TrailRestoresRunning--;
419 7      TrailRestoresLeft--;
420 6  }

421 7  if (0 != DeactivateTrailQueue(TrailID, &temp_status))
422 8  {
423 8      (void)rbe_user_error(0,
424 8          "Internal error: Cannot
425 8      deactivate trail queue for trailid '%d', cannot continue.", TrailID);
426 8  }

427 8  return -1;
428 7  }
429 7  }
430 7  }

```

```

432 7  /* This test is to determine,
433 7  * there may be another not yet started trail,
434 7  * drive available the next trail restore will be
435 7  * started.
436 7  */
437 7  if ((0 != TrailRestoresLeft) &&
438 7  (HighestActiveTrail < TrailRestoresTotal) &&
439 7  (CountDrivesInUse < CountDrivesAvailable))
440 8  {
441 8      HighestActiveTrail++;
442 7  }

443 8  if (0 != ActivateTrailQueue(HighestActiveTrail,
444 8      1,
445 8      &temp_status))
446 9  {
447 9      (void)rbe_user_error(0,
448 9          "Internal error: Cannot
449 9      activate trail queue(2) for trailid '%d', cannot continue.",
450 9      HighestActiveTrail);
451 8  }

453 8  return -1;

454 9  if (0 != SetQDrivesAcquired(
455 9      HighestActiveTrail, 1, &temp_status))
456 9  {
457 9      (void)rbe_user_error(0,
458 9          "Internal error: Cannot set
459 9      drive acquired(2) for trailid '%d', cannot continue.",
460 9      HighestActiveTrail);
461 8  }

462 8  if (0 > (temp_status = RunWorkItemRestoresForTrail(
463 8      HighestActiveTrail,
464 8      CountDrivesAvailable,
465 8      CancelRestoreTest,
466 8      &QuitFlag,
467 8      &CountDrivesInUse)))
468 9  {
469 8      return -1;
470 8  }

471 9  if (temp_status == 0)
472 9  {
473 9      /* If this Trail Had no work items we
474 9      * Should attempt to run the next trails
475 9      * work items here. This would be an internal
476 9      * error if a trail queue had no work item
477 9      * restores.
478 9      */
479 9      (void)rbe_log_stats(0,
480 9          "Internal error: Trail '%d'
481 9      restore had no work item to run(1).", HighestActiveTrail);
482 8  }

483 9  if (temp_status > 0)
484 9  {

```

```
484 9      /* IF at least on work item was started for this
485 9      */
486 9      /* then we have started a new trail.
487 9      */
488 9      TrailRestoresRunning++;
489 7      }
490 6      }
491 6      }
492 5      }
493 4      } /* end for() */
494 3
495 2      } /* else Available fds */
496 2
497 2      /*
498 2      * Terminate the loop if either
499 2      * -----
500 2      * 1) Sent the work items the quit AND
501 2      * 2) No Trail restores a running.
502 2      * OR
503 2      * 2) No more Trail restores are left.
504 2      */
505 2
506 2      if(((0 == TrailRestoresRunning) && (SentQuit)) ||
507 2      (0 == TrailRestoresLeft))
508 2      {
509 3          break;
510 3      }
511 2
512 1      } /* end while(1) */
513 1      if((0 == TrailRestoresRunning) && (SentQuit))
514 1      {
515 2          (void)rbe_log_stats(0,
516 2          "Restore was quit by user.
517 2          Work item restore quit.");
518 1      }
519 1      return 0;
520 1      }
```

```
522      /* Functions needed
523      SendRunningWorkItemsQuit();
524      InterpretWorkItemRestoreResults();
525      */
526
527      static int
528      Select(int nfd,
529      fd_set *readfds,
530      fd_set *writefds,
531      fd_set *exceptfds,
532      struct timeval *timeout)
533 1      {
534 1          int retSelect;
535 1
536 1          do
537 2          {
538 2              retSelect = select(nfd,
539 2              readfds,
540 2              writefds,
541 2              exceptfds,
542 2              timeout);
543 1
544 1          } while ((-1 == retSelect) && (EINTR == errno));
545 1          return retSelect;
546 1
547 1      }
```

```

550     eperno
551     InitiateWorkItemRestore(const int SubmitObjID,
552                             const int SubmitElemID)
553     {
554         struct auxproc AuxprocVitals;
555         eerrno_t StartupAPResults = EXIT_FAILURE;
556         time_t StartTime;
557         int TempStatus;
558         char junk_executable[1024];
559         char **junk_argv;
560         char **AP_env = NULL;
561         int SOSTatus;
562         char clientName[256] = "";
563         int clientPort;
564         int status;
565         time_t EndTime;
566         /* Lets see if there are any environment variables to set.
567          * The restore of the output variables are ignored.
568          */
569         if(E_SUCCESS != GetSOExecutionPhase(SubmitObjID,
570                                             junk_executable, 1024,
571                                             &junk_argv,
572                                             &AP_env,
573                                             &SOSTatus))
574         {
575             (void)rbe_user_error(0,
576                                 "Internal Error: Could not get environment
577                                 variables.");
578             return -1;
579         }
580         if (E_SUCCESS == GetSERcmdConnect(SubmitObjID,
581                                           SubmitElemID,
582                                           clientName, 256,
583                                           &clientPort,
584                                           &SOSTatus))
585         {
586             StartupAPResults = StartupAuxprocess(0 /* XXX */,
587                                                  &AuxprocVitals,
588                                                  AP_env,
589                                                  clientName,
590                                                  clientPort);
591         }
592         else
593         {
594             (void)rbe_user_error(0,
595                                 "Internal Error: Could not get Remote Client name
596                                 &x"
597                                 "port to connect.");
598             return -1;
599         }
600         if(E_SUCCESS != StartupAPResults)
601         {
602             /* StartupAuxprocess does its own logging. */
603             return -1;
604         }
605         /*
606          * We need to close the bulk fd. This file descriptor
607          * is not being used any more. If we do not close it
608          * here we will have a file descriptor leak because
609          * we won't be able to determine what it was when the
610          * work item completes.
611          */

```

```

612     */
613     close(AuxprocVitals.xp_fd_bulk_to_x);
614     time(&StartTime);
615     if(0 != newHandleSet(AuxprocVitals.xp_fd_to_x,
616                          AuxprocVitals.xp_fd_from_x,
617                          AuxprocVitals.xp_fd_prog_from_x,
618                          SubmitObjID,
619                          SubmitElemID,
620                          AuxprocVitals.xp_pid,
621                          StartTime,
622                          &TempStatus))
623     {
624         (void)rbe_user_error(
625             0, "Internal Error: Could not register handle set.");
626         return -1;
627     }
628     if(0 > StartWorkItemRestore(tcp,
629                                 &AuxprocVitals,
630                                 SubmitObjID,
631                                 SubmitElemID))
632     {
633         /*
634          * StartWorkItemRestore does logging if initialization fails
635          */
636         (void)rbe_user_error(
637             0, "Error in StartWorkItemRestore SubmitObjID %d,
638             SubmitElemID %d", SubmitObjID,
639             SubmitElemID);
640         /*
641          * the following code kills auxproc when recx or xcpio do not
642          * start
643          * we do not want an auxproc sitting around.
644          * If errors occur in deleteHandleSet or KillWorkItemRestore the
645          * messages are
646          * logged in those calls, plus,
647          * we already know there was an error and that
648          * is why we are doing this right now.
649          */
650         deleteHandleSet(
651             AuxprocVitals.xp_fd_from_x, EndTime, EP_RB_RECOVER_ALLFAIL, &status);
652         KillWorkItemRestore(
653             AuxprocVitals.xp_pid, AuxprocVitals.xp_fd_to_x);
654         return -1;
655     }
656     return 0;
657     /* InitiateWorkItemRestore() */
658 }
659

```

```
663  /*
664  *      Interpretate return.
665  *      Drain progress.
666  *      Send final Progress for work item.
667  *      Delete the handle set.
668  */
669
670 static int
671 HandleWorkItemRestoreResults (int FromFD,
672                               int *TrailID,
673                               wi_restore_results *results)
674 {
675     int ret = 0;
676     int retries = 0;
677     int GetAuxprocResultsStatus;
678     int TempStatus;
679     int DrainResult;
680     int DrainedFD;
681     int wiCount;
682     int AuxProcId;
683     int ToFD, getFromFD, ProgressFD;
684     time_t EndTime;
685     unsigned long jobstat;
686     int timeout = 3; /* Lets try 3 seconds */
687     boolean_t fromFDHangUp = FALSE;
688
689     ToFD = getFromFD = ProgressFD = -1;
690
691     while (! (fromFDHangUp))
692     {
693         GetAuxprocResultsStatus = GetAuxprocResults (FromFD, results);
694
695         if (-1 == GetAuxprocResultsStatus)
696         {
697             /* GetAuxprocResults() does its own logging */
698             (void)rbe_user_error(0, " Error in GetAuxprocResults");
699             return -1;
700         }
701         if (0 == GetAuxprocResultsStatus)
702         {
703             if (test_fd_hup(FromFD) == 1)
704             {
705                 fromFDHangUp = TRUE;
706             }
707         }
708
709         /* The remote result are not always going to
710          * be set. For example if the remote command
711          * is not started correctly.
712          */
713         if (results->local_exit_set == TRUE)
714         {
715             break;
716         }
717         else
718         {
719             sleep(1);
720             test_fd(FromFD);
721             continue;
722         }
723     }
724 }
725
```

```
727     time (&EndTime);
728
729     if (0 != PushDrainRequest (FromFD, &TempStatus))
730     {
731         (void)rbe_user_error(0,
732                               "Internal error: Could not push drain
733                               request, cannot continue.");
734         return -1;
735     }
736     /* Lets give the progress thread a chance to drain keeping busy in
737      * the meanwhile.
738      */
739     if (0 != FindTrailQueueOfWI (FromFD, TrailID, &TempStatus))
740     {
741         (void)rbe_user_error(0,
742                               "Internal error: Could not find trail id for
743                               finished work item, cannot continue.");
744         return -1;
745     }
746     if (0 != DecrementRunningWI (*TrailID, &wiCount, &TempStatus))
747     {
748         (void)rbe_user_error(0,
749                               "Internal error: Could not decrement running
750                               work items for trail, cannot continue.");
751         return -1;
752     }
753     if (0 != getpid (FromFD, &AuxProcId, &TempStatus))
754     {
755         (void)rbe_user_error(0,
756                               "Internal error: Could not get auxproc pid
757                               for work item, cannot continue.");
758         return -1;
759     }
760     if (0 != getHandleSet (
761         FromFD, &ToFD, &getFromFD, &ProgressFD, &TempStatus))
762     {
763         (void)rbe_user_error(0,
764                               "Internal error: Could not get auxproc file
765                               descriptors for work item, cannot continue.");
766         return -1;
767     }
768     while (0 != (ret = PopDrainResult (timeout,
769         &DrainResult,
770         &DrainedFD,
771         &TempStatus)) && retries < 3)
772     {
773         if (ret != 0)
774         {
775             "Internal error: mismatch on from file
776             descriptors for work item, cannot continue.";
777             return -1;
778         }
779         while (0 != (ret = PopDrainResult (timeout,
780             &DrainResult,
781             &DrainedFD,
782             &TempStatus)) && retries < 3)
783         {
784             if (ret != 0)
785             {

```

Page 15 of 144	HandleWorkItemRestoreResults	Thu Jan 03 12:25:21 2008
786 2	(void)rbe_user_error(0,	
787 2	"Internal error: Could not pop drain results,	
	cannot continue.");	
789 2	return -1;	
790 1	}	
792 1	/*	
793 1	Send final Progress for work item. XXX	
794 1	*/	
796 1	/* Translate the local and remote error statuses	
797 1	* to an errno value:	
798 1	*/	
800 1	if(0 != results->local_exit_status) /* use local error, if any */	
801 1	switch (results->local_exit_status)	
802 2	{	
803 2	case XG_EXIT_ALLFAIL:	
804 2	jobstat = EP_RB_RECOVER_ALLFAIL;	
805 2	break;	
806 2	case XG_EXIT_MANYFAIL:	
807 2	jobstat = EP_RB_RECOVER_MANYFAIL;	
808 2	break;	
809 2	case XG_EXIT_FEWFAIL:	
810 2	jobstat = EP_RB_RECOVER_FEWFAIL;	
811 2	break;	
812 2	case SPEXIT_REMOTE_STDBERR_PROTOCOL:	
813 2	case SPEXIT_REMOTE_STDBERR_FAIL:	
814 2	jobstat = EP_RB_RECOVER_CLIENT_STDBERR_FAIL;	
815 2	break;	
816 2	case XG_EXIT_STOPPED:	/* treat like signal */
817 2	default: /* check for signal termination vs all generic failures */	
818 2	if (XG_EXIT_SIGBASE < results->local_exit_status	
819 2	XG_EXIT_STOPPED == results->local_exit_status)	
820 3	{ /* killed by signal or stopped; separate error for sigpipe */	
821 3	if (XG_EXIT_SIGBASE + SIGPIPE == results->local_exit_status)	
822 3	jobstat = EP_RB_RECOVER_SERVER_SIGPIPE;	
823 3	else	
824 3	jobstat = EP_RB_RECOVER_SERVER_SIGNAL;	
825 2	}	
826 2	else	
827 3	{ /* generic server failure, unless client failed too */	
828 3	jobstat = EP_RB_RECOVER_SERVERFAIL;	
829 3	if (0 != results->remote_exit_status)	
830 3	jobstat = EP_RB_RECOVER_BOTHFAIL;	
831 2	}	
832 1	else if(0 != results->remote_exit_status)	
833 1	jobstat = EP_RB_RECOVER_CLIENTFAIL;	
834 1	else	
835 1	jobstat = E_SUCCESS;	
836 1	}	
838 1	if((0 != results->remote_exit_status)	
839 1	(0 != results->local_exit_status))	
840 2	{	
841 2	int status=0;	
842 2	int rc=0;	
843 2	char *templateName=NULL;	
844 2	char *wname=NULL;	
845 2	char *trailsetName=NULL;	
847 2	rc = getHandleSetInformation(FromFD,	
848 2	&templateName,	
849 2	&wname,	

Page 16 of 144	HandleWorkItemRestoreResults	Thu Jan 03 12:25:21 2008
850 2	&trailsetName,	
851 2	&status);	
852 2	rbe_log_stats(0, "Restore Failure of \"	
853 2	"top level object: %s, template %s.",	
854 2	STR_SURE(wname),	
855 2	STR_SURE(templateName));	
856 2	free(templateName);	
857 2	free(wname);	
858 2	free(trailsetName);	
859 1	}	
861 1	if(0 != deleteHandleSet(FromFD, EndTime, jobstat, &TempStatus))	
862 2	{	
863 2	(void)rbe_user_error(0,	"Internal error: Could not delete Handle
864 2	Set, cannot continue.");	
866 2	return -1;	
867 1	}	
869 1	if(0 != KillWorkItemRestore(AuxProcPid,	
870 1	-1 /* Hack this arg is not needed yet	cmd_to */))
871 2	{	
872 2	(void)rbe_user_error(0,	"Internal error: Could not kill finished
873 2	auxproc, cannot continue.");	
874 2	return -1;	
875 1	}	
877 1	close(ToFD);	
878 1	close(FromFD);	
879 1	close(ProgressFD);	
881 1	if(debugmode)	
882 2	{	
883 2	(void)rbe_user_error(0,	"DEBUG: HandleWorkItemRestoreResults Auxproc(
884 2	PID %d) just finished for trailid %d work items left = %d.",	
885 2	AuxProcPid,	
886 2	*TrailID,	
887 2	wiCount);	
889 2	(void)rbe_user_error(0,	"DEBUG: HandleWorkItemRestoreResults Auxproc(
890 2	PID %d) results are local: %d setp:%s remote: %d set:%s.",	
891 2	AuxProcPid,	
892 2	results -> local_exit_status,	
893 2	{	
894 2	results -> local_exit_set) ? "TRUE": "FALSE"	
895 2	, results -> remote_exit_status,	
	{	
	results -> remote_exit_set) ? "TRUE": "FALSE";	
898 1	}	
899 1	return 0;	
900	} /* End HandleWorkItemRestoreResults() */	

```

904  /*
905  * RunWorkItemRestoresForTrail()
906  *
907  * Description
908  * This function starts all the work item for the
909  * trail. For no this is set to one but concurrency
910  * will can be supported.
911  *
912  * Args:
913  * (I) TrailID -- The id for this trail.
914  * (I) CountDrivesAvailable -- the total drives available to restore.
915  * (O) QuitFlag -- indicate whether the user has quit the restore.
916  * (O) CountDrivesInUse -- The count of trails in use by restore.
917  *
918  * Return int
919  * if -1 then an error has occurred.
920  * if 0 or greater then the number of trail restores started will be
921  * returned.
922  */
923
924 static int
925 RunWorkItemRestoresForTrail(const int TrailID,
926                             const int CountDrivesAvailable,
927                             boolean_ty (*CancelRestorest)(),
928                             boolean_ty *QuitFlag,
929                             int *CountDrivesInUse)
930 {
931     int DriveAcquiredForTrail;
932     int DriveConcurrencyForTrail;
933     int submitObjID;
934     int submitElementID;
935     int popResults = 0;
936     int temp_status;
937     int CountOfWorkItemRestoresStarted = 0;
938     int wiCount;
939
940     while(1)
941     {
942         (*CountDrivesInUse)++;
943
944         if((0 != (popResults = PopWIFromTrailQueue(TrailID,
945                                                     &submitObjID,
946                                                     &submitElementID,
947                                                     &temp_status))) &&
948            (SCHED_NO_MORE_JOBS != temp_status))
949         {
950             (void)rbe_user_error(0,
951                                 "Internal error: Cannot pop work item off
952                                 trail queue, cannot continue.");
953             return -1;
954         }
955
956         if((-1 == popResults) && (SCHED_NO_MORE_JOBS == temp_status))
957         {
958             return CountOfWorkItemRestoresStarted;
959         }
960
961         temp_status = InitiateWorkItemRestore(
962             submitObjID, submitElementID);
963     }
964 }

```

```

965 3      {
966 3          /* InitiateWorkItemRestore() does its own logging */
967 3          (void)rbe_user_error(0, "Error in InitiateWorkItemRestore,"
968 3              "submitObjID %d, submitElementID %d", submitObjID,
969 3                  submitElementID);
970 2      }
971 2      return -1;
972 2  }
973 2      if( 0 != IncrementRunningWT(TrailID, &wCount, &temp_status))
974 2      {
975 2          (void)rbe_user_error(0,
976 2              "Internal error: Could not increment
977 2              running work items for trail, cannot continue.");
978 2      }
979 2      CountOfWorkItemRestoreStarted++;
980 2  }
981 2      if( 0 != GetTQDrivesAcquired(TrailID,
982 2          &DriveAcquiredForTrail,
983 2          &temp_status))
984 2      {
985 2          (void)rbe_user_error(0,
986 2              "Internal error: Cannot get drives
987 2              acquired, cannot continue.");
988 2      }
989 2      return -1;
990 2  }
991 2      if( 0 != GetTQDriveConcurrency(TrailID,
992 2          &DriveConcurrencyForTrail,
993 2          &temp_status))
994 2      {
995 2          (void)rbe_user_error(0,
996 2              "Internal error: Cannot get drive
997 2              concurrency, cannot continue.");
998 2      }
999 2      *QuitFlag = CancelRestoreTest();
1000 2  }
1001 2      if( (DriveAcquiredForTrail < DriveConcurrencyForTrail) &&
1002 2          ((CountDrivesInUse) < CountDrivesAvailable) &&
1003 2          (FALSE == *QuitFlag))
1004 2      {
1005 2          continue;
1006 2      }
1007 2      else
1008 2      {
1009 2          break;
1010 2      }
1011 2  }
1012 2  }
1013 2  }
1014 2  }
1015 2  }
1016 2  }
1017 2  }
1018 2  }
1019 2  }
1020 2  }
1021 2  }
1022 2  }
1023 2  }
1024 2  }
1025 2  }
1026 2  }
1027 2  }
1028 2  }
1029 2  }
1030 2  }
1031 2  }
1032 2  }
1033 2  }
1034 2  }
1035 2  }
1036 2  }
1037 2  }
1038 2  }
1039 2  }
1040 2  }
1041 2  }
1042 2  }
1043 2  }
1044 2  }
1045 2  }
1046 2  }
1047 2  }
1048 2  }
1049 2  }
1050 2  }
1051 2  }
1052 2  }
1053 2  }
1054 2  }
1055 2  }
1056 2  }
1057 2  }
1058 2  }
1059 2  }
1060 2  }
1061 2  }
1062 2  }
1063 2  }
1064 2  }
1065 2  }
1066 2  }
1067 2  }
1068 2  }
1069 2  }
1070 2  }
1071 2  }
1072 2  }
1073 2  }
1074 2  }
1075 2  }
1076 2  }
1077 2  }
1078 2  }
1079 2  }
1080 2  }
1081 2  }
1082 2  }
1083 2  }
1084 2  }
1085 2  }
1086 2  }
1087 2  }
1088 2  }
1089 2  }
1090 2  }
1091 2  }
1092 2  }
1093 2  }
1094 2  }
1095 2  }
1096 2  }
1097 2  }
1098 2  }
1099 2  }
1100 2  }
1101 2  }
1102 2  }
1103 2  }
1104 2  }
1105 2  }
1106 2  }
1107 2  }
1108 2  }
1109 2  }
1110 2  }
1111 2  }
1112 2  }
1113 2  }
1114 2  }
1115 2  }
1116 2  }
1117 2  }
1118 2  }
1119 2  }
1120 2  }
1121 2  }
1122 2  }
1123 2  }
1124 2  }
1125 2  }
1126 2  }
1127 2  }
1128 2  }
1129 2  }
1130 2  }
1131 2  }
1132 2  }
1133 2  }
1134 2  }
1135 2  }
1136 2  }
1137 2  }
1138 2  }
1139 2  }
1140 2  }
1141 2  }
1142 2  }
1143 2  }
1144 2  }
1145 2  }
1146 2  }
1147 2  }
1148 2  }
1149 2  }
1150 2  }
1151 2  }
1152 2  }
1153 2  }
1154 2  }
1155 2  }
1156 2  }
1157 2  }
1158 2  }
1159 2  }
1160 2  }
1161 2  }
1162 2  }
1163 2  }
1164 2  }
1165 2  }
1166 2  }
1167 2  }
1168 2  }
1169 2  }
1170 2  }
1171 2  }
1172 2  }
1173 2  }
1174 2  }
1175 2  }
1176 2  }
1177 2  }
1178 2  }
1179 2  }
1180 2  }
1181 2  }
1182 2  }
1183 2  }
1184 2  }
1185 2  }
1186 2  }
1187 2  }
1188 2  }
1189 2  }
1190 2  }
1191 2  }
1192 2  }
1193 2  }
1194 2  }
1195 2  }
1196 2  }
1197 2  }
1198 2  }
1199 2  }
1200 2  }
1201 2  }
1202 2  }
1203 2  }
1204 2  }
1205 2  }
1206 2  }
1207 2  }
1208 2  }
1209 2  }
1210 2  }
1211 2  }
1212 2  }
1213 2  }
1214 2  }
1215 2  }
1216 2  }
1217 2  }
1218 2  }
1219 2  }
1220 2  }
1221 2  }
1222 2  }
1223 2  }
1224 2  }
1225 2  }
1226 2  }
1227 2  }
1228 2  }
1229 2  }
1230 2  }
1231 2  }
1232 2  }
1233 2  }
1234 2  }
1235 2  }
1236 2  }
1237 2  }
1238 2  }
1239 2  }
1240 2  }
1241 2  }
1242 2  }
1243 2  }
1244 2  }
1245 2  }
1246 2  }
1247 2  }
1248 2  }
1249 2  }
1250 2  }
1251 2  }
1252 2  }
1253 2  }
1254 2  }
1255 2  }
1256 2  }
1257 2  }
1258 2  }
1259 2  }
1260 2  }
1261 2  }
1262 2  }
1263 2  }
1264 2  }
1265 2  }
1266 2  }
1267 2  }
1268 2  }
1269 2  }
1270 2  }
1271 2  }
1272 2  }
1273 2  }
1274 2  }
1275 2  }
1276 2  }
1277 2  }
1278 2  }
1279 2  }
1280 2  }
1281 2  }
1282 2  }
1283 2  }
1284 2  }
1285 2  }
1286 2  }
1287 2  }
1288 2  }
1289 2  }
1290 2  }
1291 2  }
1292 2  }
1293 2  }
1294 2  }
1295 2  }
1296 2  }
1297 2  }
1298 2  }
1299 2  }
1300 2  }
1301 2  }
1302 2  }
1303 2  }
1304 2  }
1305 2  }
1306 2  }
1307 2  }
1308 2  }
1309 2  }
1310 2  }
1311 2  }
1312 2  }
1313 2  }
1314 2  }
1315 2  }
1316 2  }
1317 2  }
1318 2  }
1319 2  }
1320 2  }
1321 2  }
1322 2  }
1323 2  }
1324 2  }
1325 2  }
1326 2  }
1327 2  }
1328 2  }
1329 2  }
1330 2  }
1331 2  }
1332 2  }
1333 2  }
1334 2  }
1335 2  }
1336 2  }
1337 2  }
1338 2  }
1339 2  }
1340 2  }
1341 2  }
1342 2  }
1343 2  }
1344 2  }
1345 2  }
1346 2  }
1347 2  }
1348 2  }
1349 2  }
1350 2  }
1351 2  }
1352 2  }
1353 2  }
1354 2  }
1355 2  }
1356 2  }
1357 2  }
1358 2  }
1359 2  }
1360 2  }
1361 2  }
1362 2  }
1363 2  }
1364 2  }
1365 2  }
1366 2  }
1367 2  }
1368 2  }
1369 2  }
1370 2  }
1371 2  }
1372 2  }
1373 2  }
1374 2  }
1375 2  }
1376 2  }
1377 2  }
1378 2  }
1379 2  }
1380 2  }
1381 2  }
1382 2  }
1383 2  }
1384 2  }
1385 2  }
1386 2  }
1387 2  }
1388 2  }
1389 2  }
1390 2  }
1391 2  }
1392 2  }
1393 2  }
1394 2  }
1395 2  }
139
```

```

1018 /* Stub */
1019 static int DetermineGlobalDriveUse()
1020 {
1021     /* Limiting to MAXINT == not limiting... Need resource management
1022     * to do this properly.
1023     NOTE: This should now work like eb_dc_restore does.
1024     */
1025     return MAXINT;
}

```

```

1028 static int
1029 SendRunningWorkItemsQuit()
1030 {
1031     int *APlist;
1032     int count;
1033     int status;
1034     int index;
1036     if(0 != getPIDList(&count, &APlist, &status))
1037     {
1038         (void)rbe_user_error(0,
1039                             "Internal error: Cannot get auxproc pid list,
1040                             cannot continue.");
1041         return -1;
1042     }
1043     for(index = 0; index < count; index++)
1044     {
1045         QuitWorkItemRestore(APlist[index]);
1046     }
1047     return 0;
1048 }

```

```
1050  /*
1051  * Stub this out for now.
1052  */
1053  static int
1054  InterpretWorkItemRestoreResults(wl_restore_results *results)
1055  {
1056      return 0;
1057  }
```

```
1059  static void
1060  DebugLogFds(char *error_msg,
1061             fd_set *fds)
1062  {
1063      int index, fd_count = 0;
1064      char buffer[4096];
1065      char *bufptr = (char*)buffer;
1066
1067      for(index=0;
1068          index < 1024;
1069          index++)
1070      {
1071          if( FD_ISSET(index, fds))
1072          {
1073              int size = 0;
1074              size = sprintf(bufptr, "%d, ", index);
1075              bufptr += size;
1076              fd_count++;
1077          }
1078      }
1079      rbe_log_stats(0, "%s fd_count:: %d :: (%s)\n",
1080                  error_msg, fd_count, buffer);
1081  }
1082  }
```



```

1085 static int
1086 test_fd(int fd)
1087 {
1088     fd_set read_fd;
1089     int ret_select;
1090     struct timeval timeout = {0, 0};
1092     FD_ZERO(&read_fd);
1094     FD_SET(fd, &read_fd);
1096     do
1097     {
1098         ret_select = select(fd + 1, &read_fd, NULL, NULL, &timeout);
1099     } while((-1 == ret_select) && (EINTR == errno));
1100     return ret_select;
1102 }
1104

```

```

1106 /*
1107  * test_fd_hup()
1108  * Description: Test the supplied file descriptor to see if
1109  * it has had the hang up condition.
1110  *
1111  * Args:
1112  * Input fd -- the file descriptor to check for the hang up condition.
1113  *
1114  * Returns:
1115  * 1 for HUP event received on fd.
1116  * 0 No HUP event received on fd.
1117  * -1 errno set.
1118  *
1119  */
1120 static int
1121 test_fd_hup(int fd)
1122 {
1123     struct pollfd fds;
1124     int ret_poll;
1125
1126     if(fd < 0)
1127     {
1128         errno = EINVAL;
1129         return -1;
1130     }
1131
1132     fds.fd = fd;
1133     fds.events = POLLIN;
1134     fds.revents = 0; /* initialize */
1135
1136     do
1137     {
1138         ret_poll = poll(&fds, 1, 0);
1139     } while((-1 == ret_poll) && (EINVAL == errno));
1140
1141     if(-1 == ret_poll)
1142     {
1143         return -1;
1144     }
1145
1146     if(POLLHUP & fds.revents)
1147     {
1148         return 1;
1149     }
1150     else
1151     {
1152         return 0;
1153     }
1154 }
1155
1156 } /* end test_fd_hup() */
1157

```



```
1 /*****
2 **
3 ** File Name: RSLauxmgr.c
4 **
5 ** Copyright (c) 1998,1999 by EMC Corporation.
6 **
7 ** Purpose:
8 ** -----
9 ** The intent of the contents of this file is to implement the
10 ** functions the control execution of work item restores Or the
11 ** running of auxproc.
12 **
13 **
14 ** Library.
15 **
16 ** These functions are provided to allow:
17 ** - The pipe, fork, duping closing and exec of auxproc.
18 ** - The starting of work item restores.
19 ** - The quitting of work item restores.
20 ** - The getting results of work item restores.
21 **
22 **
23 ** The following functions comprise restoral management:
24 **
25 **
26 **
27 **
28 **
29 **
30 ** Compile-Time Options:
31 ** This section must list any compile time definitions
32 ** which will affect this header.
33 **
34 *****/
35
36 /*
37 * Feature test switches.
38 * Standard defines required to turn on OS features go here.
39 *
40 * The following is required for code that uses POSIX API's.
41 * Remove for non-POSIX, non-portable code.
42 */
43
44 #define _POSIX_SOURCE 1
45
46 /*
47 * System headers.
48 */
49
50 #include <sys/wait.h>
51 #include <sys/types.h>
52 #include <unistd.h>
53 #include <string.h>
54 #include <stdlib.h>
55
56 /*
57 * Epoch headers.
58 */
59
60 #include <eb/eb_port.h>
61 #include <eb/rb_log.h>
62 #include <ebutil/eb_normalize.h>
63 #include <ebutil/ebutil.h>
64 #include <ebreport/ebv1.h>
```

```
65 /*
66 * Local headers
67 */
68
69 #include <RSLinterns.h>
70 #include <RSLrbmain.h>
71 #include <restore/EDMRSubmittapi.h>
72 #include <EDMfork.h>
73 #include <RSLauxsupp.h>
74
75 #define SUBMIT_FIELD_MAX 2048
76
77 extern int putenv(const char *string);
78
79 extern char *strsignal(int sig);
80
81 extern int ChildDone(int child_pid, int *child_result);
82
83 static void
84 reset_recovery_privileges(struct recover_context *rcx,
85                          int changed);
86
87 static void
88 set_recovery_privileges(struct recover_context *rcx,
89                        int *changed);
90
91 static char *
92 generate_rcmddpath(int SubmitObjectID, int SubmitEID);
93
94 /*
95 * Startupauxprocess()
96 *
97 * Description:
98 * This function pipe, fork, & exec auxproc. After which it tests
99 * its the just started auxproc with the Ping command.
100 *
101 * processing is it closes the fds that are not associated with
102 * but are inherited from the parent (
103 * restore engine). If any environment
104 * variable need to be set for auxproc, they are set.
105 *
106 * Args:
107 * (Inp) debugmode -- int 1 for debug and 0 for no debug.
108 * (Out) struct auxproc *xp -- preallocated structure to return vital
109 * info
110 * (char) **auxproc_envp -- environment variables to be appended to
111 * auxproc's
112 * environment.
113 * The rest of the environment is
114 * inherited from auxproc.
115 * Format is "ENV=value\0"
116 * char *socketClientNm -- used for the client initiated restore,
117 * this is the
118 * client name
119 * int clientSocketPort -- Used to identify the port number on which
120 * to
121 * contact the client
122 *
123 * Notes:
124 * auxproc_envp can be NULL indicating no environment to append to
125 * auxproc.
126 * This is a char ** which last char * should be NULL.
127 *
128 * Inherited from:
129
```

```

122  * static eerrno_tly setup_aux_processes(struct recover_context *rcx)
123  /*
124
125  eerrno_tly
126  StartupAuxprocess(int
127      struct          debugmode,
128      auxproc *xp,    **auxproc_envp,
129      char            *socketClientNm,
130      int              clientSocketPort)
131  {
132
133      /*
134      * NOTES:
135      * 1) Do I really want to be reliant on the recover_context struct.
136      * 2) I need to fork() exec() auxproc.
137      * 3) Is fork1 really going to work.
138      */
139
140  #define RFD 0 /* in a pipe, fd 0 is the read descriptor */
141  #define WFD 1 /* and fd 1 is the write descriptor */
142
143  int save_errno;
144  int fd;
145  char *resultsbuf = NULL;
146  char *auxproc_pathname = AUXPROC_PATHNAME;
147  char *auxproc_executable = AUXPROCNAME;
148
149  int ping_status;
150  int auxproc_index = 0; /* alton Remove */
151  int cmd_pipe_to[2];
152  int cmd_pipe_from[2];
153  int bulk_pipe_to[2];
154  int prog_pipe_from[2];
155
156  if (pipe(cmd_pipe_to) == -1 ||
157      pipe(cmd_pipe_from) == -1 ||
158      pipe(bulk_pipe_to) == -1 ||
159      pipe(prog_pipe_from) == -1)
160  {
161      save_errno = errno;
162      rbe_log_stats(RBRECOVER_MKERR(errno), "pipe() failed");
163      return(RBRECOVER_MKERR(save_errno));
164  }
165
166  /* This below appends environment variables to
167  * the environment that auxproc inherits from the
168  * restore engine.
169  */
170
171  if(NULL != auxproc_envp)
172  {
173      int index;
174      for(index=0; NULL != auxproc_envp[index]; index++)
175      {
176          if(0 != putenv(auxproc_envp[index]))
177          {
178              rbe_log_stats(RBRECOVER_MKERR(errno),
179                  "Unable to set auxproc environment %s,
180                      for auxproc PID %d.",
181                      auxproc_envp[index], getpid());
182          }
183      }
184      _exit(1); /* We are the child */
185  }
186  }

```

```

188  {
189      switch (xp->xp_pid == EDMfork1())
190      {
191          case -1: /* Error */
192              save_errno = errno;
193              rbe_log_stats(RBRECOVER_MKERR(errno), "fork() failed");
194              return(RBRECOVER_MKERR(save_errno));
195          case 0: /* child */
196              {
197                  char procnum_str[32];
198                  char r_fd_str[32];
199                  char w_fd_str[32];
200                  char w_bulk_str[32];
201                  char r_bulk_fd_str[32];
202                  char w_prog_fd_str[32];
203                  char dbgmodestr[32];
204                  char socket_host_str[256];
205
206                  /* Not sure what this should be */
207                  char socket_port_str[32];
208                  int ret_exec = 0;
209                  socket_port_str[0] = '\0';
210
211                  (void) sprintf(procnum_str, "%d", auxproc_index);
212                  (void) sprintf(r_fd_str, "%d", cmd_pipe_to[RFD]);
213                  (void) sprintf(w_fd_str, "%d", cmd_pipe_from[WFD]);
214                  (void) sprintf(r_bulk_str, "%d", bulk_pipe_to[RFD]);
215                  (void) sprintf(w_prog_fd_str, "%d", prog_pipe_from[WFD]);
216                  (void) sprintf(dbgmodestr, "%d", debugmode);
217                  if (NULL != socketClientNm)
218                  {
219                      (void) sprintf(socket_host_str, "%s", socketClientNm);
220                      (void) sprintf(socket_port_str, "%d", clientSocketPort);
221                  }
222                  else
223                  {
224                      socket_host_str[0] = 0;
225                      socket_port_str[0] = 0;
226                  }
227                  ret_exec = execvp(
228                      Auxproc_pathname, /* prog to execute */
229                      Auxproc_executable, /* argv 0 */
230                      procnum_str,
231                      r_fd_str,
232                      w_fd_str,
233                      r_bulk_fd_str,
234                      w_prog_fd_str,
235                      dbgmodestr,
236                      socket_host_str,
237                      socket_port_str,
238                      (char *)0);
239
240                  rbe_log_stats(RBRECOVER_MKERR(errno),
241                      "Unable to exec %s, for %s PID %d.",
242                      AUXPROCNAME,
243                      getpid());
244                  _exit(1);
245              }
246          default: /* parent */
247              {
248                  /*
249

```

```

251 2 /*
252 2 * The parent has no need for the fd
253 2 * used to write *to* the parent, nor
254 2 * the fds used to read *from* the parent.
255 2 * If these are not close EPIPE and SIGPIPE
256 2 * may be missed by the writer when the
257 2 * intended reader dies.
258 2 */
260 2 (void)close(cmd_pipe_to[RFD]);
261 2 (void)close(bulk_pipe_to[RFD]);
262 2 (void)close(cmd_pipe_from[WFD]);
263 2 (void)close(prog_pipe_from[WFD]);
265 2 /*
266 2 * but does want to save the other fds
267 2 */
269 2 xp->xp_fd_from_x = cmd_pipe_from[RFD];
270 2 xp->xp_fd_to_x = cmd_pipe_to[WFD];
271 2 xp->xp_fd_bulk_to_x = bulk_pipe_to[WFD];
272 2 xp->xp_fd_prog_from_x = prog_pipe_from[RFD];
273 2 /*
274 2 * In debugmode, exec the separate-process
275 2 * version of the auxprocs (ptrace issue)
276 2 */
278 2 if (debugmode)
279 2 {
280 3 auxcmdpacket(xp->xp_fd_to_x, 'X', 0, "");
281 2 }
282 1 }
283 1 #define PING_TEST_STR "abc"
284 1 #define PING_TEST_STR_SIZE 4 /* include the '\0' */
286 1 fd = xp->xp_fd_to_x;
287 1 auxcmdpacket(fd, 'P', PING_TEST_STR_SIZE, PING_TEST_STR);
289 1 fd = xp->xp_fd_from_x;
291 1 ping_status = auxresults(fd, 'P', 0, &resultbuf);
293 1 if ((-1 == ping_status) ||
294 1 (NULL == resultbuf) ||
295 1 (strcmp(resultbuf, PING_TEST_STR) != 0))
296 2 {
297 2 rec_api_log_csm(SUB_CSM_NO_PING_AUXPROC, NULL);
298 2 rbe_log_stats(0, "%s ping start-up",
299 2 AUXPROCNAME);
300 2 return(EP_RB_RECOVER_AUXPROC_DIED);
301 1 }
303 1 free(resultbuf);
305 1 return(E_SUCCESS);
307 1 #undef PING_TEST_STR
308 1 #undef PING_TEST_STR_SIZE
309 1 #undef RFD
310 1 #undef WFD
311 1 #undef MAX_FD
312 1 } /* end of setup_aux_process() */

```

```

315 #define MAX_SUBMIT_FIELD 2048
318 /*****
319 * start_cpigen()
320 *
321 * This function initiates a work item restore. It first determines
322 * the size of the restore command to send to auxproc, malloc's the
323 * memory and creates the restore command and sends it to auxproc.
324 * Auxproc will start the rcmd (if necessary) and start xcpigen.
325 * Auxproc will send the initialization reply which is read by this
326 * function. The results are for initialization.
327 *
328 * Args:
329 * (I) rcx -- Recover Context struct (limited use)
330 * (I) xp -- auxproc struct pointer.
331 * (I) submitObjectID -- identifies what and how to run restore.
332 * (I) submitElemID -- identifies what and how to run witem
333 * restore.
334 * (I) auxprocnum -- auxproc number may become obsolete.
335 * (I) *rcmdinfo[4] -- the rcmdinfo must be included.
336 * (O) results -- Of the work item restore initialization.
337 * (O) err_str -- Of work item initialization failures.
338 *
339 * Returns:
340 * xcpigen's pid for success, -1 for failure.
341 *
342 * NOTES:
343 * rcx -- The recover context structure is carefully used below.
344 * The rcx structure should be used only to get the global values
345 * like the xcpigen executable name and the config structure.
346 *
347 * Submit Object should be used to determine user id,
348 * admin privileges, and other values that would not vary
349 * for a potentially multi work item restore.
350 *
351 * Submit Element should be used to determine anything that
352 * could be potentially unique for a work item restore.
353 *
354 * rcmdinfo is the command line for the remote command.
355 * *****/
356 int
357 start_cpigen(struct recover_context *rcx,
358 struct auxproc *xp,
359 int submitObjectID,
360 int submitElemID,
361 int auxprocnum,
362 char *rcmdinfo[4],
363 rcmd_pkt0_info *results,
364 char **err_str)
365 {
366 1 char *auxproc_databuf;
367 1 size_t data_len = 0;
368 1 char *p;
369 1 int i;
370 1 int resfd;
371 1 /* For Initialization results */
372 1 int pid = -1;
373 1 rcmd_pkt0_info *pkt0p;
374 1 rcmd_pkt0_info *pkt0p_buffer;
375 1 char *resultbufptr = NULL;
376 1 char *errstr = "";
377 1 struct mark_summary submit_summary;

```

Page 35 of 144	start_cpiogen	Thu Jan 03 12:25:21 2008	Page 36 of 144	start_cpiogen	Thu Jan 03 12:25:21 2008
<pre>379 1 /* Args for xcpiogen */ 380 1 char *xcpiogen_argv0; 381 1 int xcpiogen_argc; 382 1 char *submit_file_flag = "-S"; 383 1 char *outputfd_fmt = "-f%d"; 384 1 char *progress_report_flag = "-p"; 385 1 u_hyper total_bytes; 386 1 char *total_bytes_stringP; 387 1 char *total_bytes_stringB; 388 1 char *bufsize = NULL; 389 1 char bufsize_buffer[20]; 390 1 391 1 RBC_WORKGROUP *pg; 392 1 RBC_WORKITEM *pi; 393 1 394 1 /* temporary variable for the submit object / element fields. */ 395 1 char temp_effective_uidname[MAX_SUBMIT_FIELD]; 396 1 char temp_socket_host[MAX_SUBMIT_FIELD]; 397 1 char temp_workitem_name[MAX_SUBMIT_FIELD]; 398 1 char temp_submit_file[MAX_SUBMIT_FIELD]; 399 1 int temp_socket_port; 400 1 int GetSEStatus = 0; 401 1 int GetSEStatus = 0; 402 1 403 1 if(GetSEWorkItemName(submitObjectID, submitElemID, 404 1 temp_workitem_name, MAX_SUBMIT_FIELD, 405 1 &GetSEStatus) != 0) 406 1 { 407 1 rbe_log_stats(0, "Unable to get work item name."); 408 1 return -1; 409 1 } 410 1 411 1 if(GetSEEffectiveUserName(submitObjectID, 412 1 temp_effective_uidname, 413 1 MAX_SUBMIT_FIELD, &GetSEStatus) != 0) 414 1 { 415 1 rbe_log_stats(0, "Unable to get user name."); 416 1 return -1; 417 1 } 418 1 419 1 if(GetSECmdConnect(submitObjectID, submitElemID, 420 1 temp_socket_host, MAX_SUBMIT_FIELD, 421 1 &temp_socket_port, &GetSEStatus) != 0) 422 1 { 423 1 rbe_log_stats(0, "Unable to get socket host/port pair."); 424 1 return -1; 425 1 } 426 1 427 1 if(GetSESubmitFile(submitObjectID, submitElemID, 428 1 temp_submit_file, MAX_SUBMIT_FIELD, 429 1 &GetSEStatus) != 0) 430 1 { 431 1 rbe_log_stats(0, "Unable to get submit file."); 432 1 return -1; 433 1 } 434 1 435 1 /* 436 1 * +1 for '\0' characters 437 1 */ 438 1 data_len += strlen(rcmdinfo[0]) + 1; 439 1 data_len += strlen(rcmdinfo[1]) + 1; 440 1 data_len += strlen(rcmdinfo[2]) + 1; 441 1 442 1 /* rcmd_hostname */ 443 1 /* rcmd_locator */ 444 1 /* rcmd_remuser */ 445 1 446 1 data_len += strlen(rcmdinfo[3]) + 1; 447 1 448 1 data_len += sizeof (int); 449 1 data_len += sizeof (int); 450 1 data_len += strlen(rcx -> rc_cpiogen_executable) + 1; 451 1 data_len += sizeof (int); 452 1 453 1 /* 454 1 * xcpiogen arguments. Pass traditional argv[0] 455 1 * plus the output file descriptor number, 456 1 * plus -p (progress report mode) and -B with its 457 1 * argument (total bytes to be processed). 458 1 */ 459 1 460 1 xcpiogen_argc = 7; 461 1 462 1 xcpiogen_argv0 = strchr(rcx -> rc_cpiogen_executable, '/'); 463 1 if (NULL != xcpiogen_argv0) 464 1 { 465 1 ++xcpiogen_argv0; 466 1 } 467 1 else 468 1 { 469 1 xcpiogen_argv0 = rcx -> rc_cpiogen_executable; 470 1 } 471 1 data_len += strlen(xcpiogen_argv0) + 1; 472 1 data_len += strlen(submit_file_flag) + 1; 473 1 data_len += strlen(temp_submit_file) + 1; 474 1 data_len += strlen(outputfd_fmt) + 1; 475 1 data_len += strlen(progress_report_flag) + 1; 476 1 477 1 /* 478 1 * Get the current total number of bytes to be processed from 479 1 * the mark summary u_hyper so that we can then convert it to 480 1 * a decimal string to be passed to xcpiogen(). 481 1 */ 482 1 if(0 != GetSEMarkedSummary(submitObjectID, 483 1 submitElemID, 484 1 &submit_summary, 485 1 &GetSEStatus)) 486 1 { 487 1 /* This is not a critical error. This may cause progress 488 1 * reporting problems! 489 1 */ 490 1 491 1 rbe_log_stats(0, "Unable to set submit size for xcpiogen."); 492 1 total_bytes = u_l_to_uh(0); 493 1 } 494 1 else 495 1 { 496 1 total_bytes = submit_summary.len_mkd_files; 497 1 } 498 1 499 1 total_bytes_stringP = u_hyper_to_decimal(total_bytes); 500 1 501 1 /* 502 1 * Add the size of the "total bytes" flag and value string 503 1</pre>	<pre>444 1 data_len += strlen(rcmdinfo[3]) + 1; 445 1 446 1 data_len += strlen(temp_effective_uidname) + 1; 447 1 448 1 data_len += sizeof (int); 449 1 data_len += sizeof (int); 450 1 data_len += strlen(rcx -> rc_cpiogen_executable) + 1; 451 1 data_len += sizeof (int); 452 1 453 1 /* 454 1 * xcpiogen arguments. Pass traditional argv[0] 455 1 * plus the output file descriptor number, 456 1 * plus -p (progress report mode) and -B with its 457 1 * argument (total bytes to be processed). 458 1 */ 459 1 460 1 xcpiogen_argc = 7; 461 1 462 1 xcpiogen_argv0 = strchr(rcx -> rc_cpiogen_executable, '/'); 463 1 if (NULL != xcpiogen_argv0) 464 1 { 465 1 ++xcpiogen_argv0; 466 1 } 467 1 else 468 1 { 469 1 xcpiogen_argv0 = rcx -> rc_cpiogen_executable; 470 1 } 471 1 data_len += strlen(xcpiogen_argv0) + 1; 472 1 data_len += strlen(submit_file_flag) + 1; 473 1 data_len += strlen(temp_submit_file) + 1; 474 1 data_len += strlen(outputfd_fmt) + 1; 475 1 data_len += strlen(progress_report_flag) + 1; 476 1 477 1 /* 478 1 * Get the current total number of bytes to be processed from 479 1 * the mark summary u_hyper so that we can then convert it to 480 1 * a decimal string to be passed to xcpiogen(). 481 1 */ 482 1 if(0 != GetSEMarkedSummary(submitObjectID, 483 1 submitElemID, 484 1 &submit_summary, 485 1 &GetSEStatus)) 486 1 { 487 1 /* This is not a critical error. This may cause progress 488 1 * reporting problems! 489 1 */ 490 1 491 1 rbe_log_stats(0, "Unable to set submit size for xcpiogen."); 492 1 total_bytes = u_l_to_uh(0); 493 1 } 494 1 else 495 1 { 496 1 total_bytes = submit_summary.len_mkd_files; 497 1 } 498 1 499 1 total_bytes_stringP = u_hyper_to_decimal(total_bytes); 500 1 501 1 /* 502 1 * Add the size of the "total bytes" flag and value string 503 1</pre>				

Page 35 of 144

RSlauxmgr.c:7

Thu Jan 03 12:25:21 2008

Page 36 of 144

RSlauxmgr.c:8

Thu Jan 03 12:25:21 2008

```

504 1      */
506 1      data_len += strlen(total_bytes_flag) + 1; /* xcploggen argv[5] */
507 1      data_len += strlen(total_bytes_stringp) + 1; /* xcploggen argv[6] */

/*
 * Add size for db API socket info
 */
513 1      data_len += sizeof (int); /* socket port # */
514 1      data_len += strlen(temp_socket_host) + 1;

/*
 * Locate work item in config info & get length of filespec
 * %s changed break in inner loop to 'goto' to resume with
 * found item.
 */
522 1      for (pi = NULL, pg = rcx->rc_config->pgrouplist /* OK */;
523 1          NULL != pg;
524 1          pg = pg->next)
525 2      {
526 2          for (pi = pg->pwlist; NULL != pi; pi = pi->next)
527 3              if (0 == strcmp(pi->name, temp_workitem_name))
528 3                  goto stopsearch; /* %s exit both loops */
529 4
530 4      }
531 3
532 2
533 1      }

stopsearch:
535 1      if (pi != NULL)
536 1      {
537 2          data_len += strlen(pi->list)+1;
538 2
539 1      }
540 1      else
541 2      {
542 2          data_len += 1;
543 1      }

if (NULL != pi && DEFAULT_BB_BUFSIZE !=
545 1      pi->recover_server_bufsize)
546 2      {
547 2          sprintf(bufsize_buffer, "-R%d", pi->recover_server_bufsize);
548 2          bufsizep = bufsize_buffer;
549 2          data_len += strlen(bufsizep) + 1;
550 2          xcploggen_argc++; /* one more arg to xcploggen */
551 1      }

data_len += strlen(temp_workitem_name) + 1;

/*
 * Allocate memory to hold all this gunk we need to
 * shove towards our auxiliary process.
 */
555 1      auxproc_databuf = sm_malloc((unsigned)data_len);
556 1
557 1
558 1
560 1
562 1      /*
563 1      * Fill in the gunk. First, the rcmd info for the rsh part.
564 1      */
566 1      p = auxproc_databuf;

```

```

567 1      for (i = 0; i < 4; i++)
568 2      {
569 2          (void)strcpy(p, rcmdinfo[i]);
570 2          p += strlen(p)+1;
571 1      }

/*
 * The human username which the remote should operate as
 */
573 1      (void)strcpy(p, temp_effective_username);
574 1      p += strlen(p)+1;
575 1

577 1
578 1

580 1      /*
581 1      * The xcploggen-cmd-fd-info, which tells auxproc which
582 1      * arg has the sprintf format string to insert the
583 1      * actual output fd number. Remember auxproc sets up
584 1      * the connections between xcploggen and the rcmd. This
585 1      * below arg is sent to auxproc to tell auxproc which
586 1      * argument to update the format string with the outputfd.
587 1      */

589 1      i = 3; /* For argv[3] -- See outputfd_fmt below.*/
591 1      memcpy(p, &i, sizeof(int));
593 1      p += sizeof (int);

/*
 * The flags, which are always zero currently
 */
595 1      i = 0;
596 1      memcpy(p, &i, sizeof (int));
597 1      p += sizeof (int);

/*
 * The "xcploggen command" that we will run locally.
 */
603 1      (void)strcpy(p, rcx -> rc_cploggen_executable);
604 1      p += strlen(p)+1;
605 1

607 1
608 1

610 1      /*
611 1      * The argc for the "xcploggen command", and its argv vector.
612 1      */
614 1      memcpy(p, &xcploggen_argc, sizeof (int));
615 1      p += sizeof (int);

617 1      /*
618 1      * The argv vector, which is argv0, the outputfd
619 1      * thing, and the progress report mode flag.
620 1      */
622 1      (void)strcpy(p, xcploggen_argv0);
623 1      p += strlen(p)+1;

625 1      (void)strcpy(p, submit_file_flag); /* xcploggen argv[1] */
626 1      p += strlen(submit_file_flag) + 1;

628 1      (void)strcpy(p, temp_submit_file); /* xcploggen argv[2] */
629 1      p += strlen(temp_submit_file) + 1;

631 1      (void)strcpy(p, outputfd_fmt); /* xcploggen argv[3] */
632 1      p += strlen(outputfd_fmt) + 1;

```



```

634 1      (void)strcpy(p, progress_report_flag); /* xcpiogen argv[4] */
635 1      p += strlen(progress_report_flag) + 1;

637 1      if (bufsizep != NULL)
638 2      {
639 2          (void)strcpy(p, bufsizep); /* xcpiogen argv[??] */
640 2          p += strlen(bufsizep) + 1;
641 1      }

643 1      /*
644 1      * Follow this with the total bytes flag and value.
645 1      */

647 1      strcpy(p, total_bytes_flag); /* xcpiogen argv[??] */
648 1      p += strlen(total_bytes_flag) + 1;
649 1      strcpy(p, total_bytes_stringp); /* xcpiogen argv[??] */
650 1      p += strlen(total_bytes_stringp) + 1;

652 1      /*
653 1      * socket info for db API
654 1      */

656 1      (void)memcpy(p, (char *)&temp_socket_port, sizeof(int));
657 1      p += sizeof(int);

659 1      /*
660 1      * socket host name for db API
661 1      */

663 1      (void)strcpy(p, temp_socket_host);
664 1      p += strlen(p)+1;

666 1      if (NULL != pi) /* send filespec */
667 2      {
668 2          (void)strcpy(p, pi->list);
669 2          p += strlen(p)+1;
670 1      }
671 1      else
672 2      {
673 2          *p++ = 0;
674 1      }

676 1      /*
677 1      * workitem name
678 1      */

680 1      (void)strcpy(p, temp_workitem_name);
681 1      p += strlen(p)+1;

683 1      /*
684 1      * assert that our arithmetic above was done correctly
685 1      */

687 1      if ((size_t)(p - auxproc_databuf) != data_len)
688 2      {
689 2          rbe_log_stats(0, "assertion failed: cmd size miscount");
690 2          return -1;
691 1      }

693 1      /*
694 1      * Send the restore command to auxproc. Auxproc will start
695 1      * The remote command (if necessary) and xcpiogen.
696 1      */

698 1      auxcmdpacket(xp-> xp_fd, to_x,

```

```

699 1      'r', (int)data_len, auxproc_databuf);

701 1      /*
702 1      * Obtain the fork status
703 1      */

705 1      resfd = xp-> xp_fd_from_x;
707 1      i = auxresults(resfd, '0', 0, &resultsbufptr);

709 1      if (i < 0)
710 2      {
711 2          rbe_log_stats(0,
712 2              "*** Error while starting auxproc for work item
713 2                  temp_workitem_name);
714 2          null_free(resultsbufptr);
715 2          return -1;
716 1      }

718 1      /* This memory management is crap */

720 1      pkt0p = &pkt0p_buffer;
721 1      memcpy(pkt0p, resultsbufptr, sizeof(pkt0p_buffer));
722 1      memcpy(results, resultsbufptr, sizeof(pkt0p_buffer));

724 1      if ((pkt0p->msglen) > 0)
725 2      {
726 2          errstr = resultsbufptr + sizeof *pkt0p;
727 2          *err_str = esl_strdup(errstr);
728 1      }
729 1      else
730 2      {
731 2          errstr = "";
732 2          *err_str = esl_strdup("");
733 1      }

735 1      /*
736 1      * if the fork failed, the cpiogen start fails
737 1      */

739 1      if (0 != pkt0p->failcode)
740 2      {
741 2          int jnk;

744 2          if (strlen(errstr) > (size_t)0)
745 3          {
746 3              rbe_log_stats(0,
747 3                  "*** Error while starting auxproc, error %s,
748 3                      \"for work item \"%s\", errstr,
749 3                      temp_workitem_name);
751 2          }
752 2          free(resultsbufptr);
753 2          /*
754 2          * collect the useless 'r' reply packet
755 2          */

757 2          resultsbufptr = (char *)&jnk;
758 2          (void)auxresults(resfd, 'r', sizeof(int), &resultsbufptr);
759 2          return -1;
760 1      }

```

```
762 1 /*
763 1 * Caller assumes responsibility for (eventually)
764 1 * collected exit status of remote and local programs.
765 1 */
767 1 pid = pktOp->pid;
769 1 free (resultsbufptr);
771 1 return pid;
772 1 } /* end of start_cpiogen() */
```

```
775 /*
776 * 1) Remove rcx references.
777 */
779 static char *
780 make_remote_cpiogen_cmd(struct recover_context *rcx,
781 int SubmitObjID,
782 int SubmitElemID)
783 {
784 char *rcmdpath = generate_rcmdpath(SubmitObjID,
785 SubmitElemID);
786 char *minus_c = "";
787 char *minus_c_arg = "";
788 char *noclobber = "";
789 char *remdebugflag = "";
790 char *cmd;
791 unsigned len;
792 RBC_WORKITEM *work_item;
793 char *bufsizep = "";
794 char bufsize_buffer[20];
796 char temp_dirtop[SUBMIT_FIELD_MAX];
797 char temp_workitem_name[SUBMIT_FIELD_MAX];
798 OverwritePolicy temp_overwrite_policy;
799 int GetSEStatus = 0;
801 if (NULL == rcmdpath)
802 {
803 return NULL;
804 }
805 temp_overwrite_policy = GetSEDestOverWritePolicy(SubmitObjID,
806 SubmitElemID,
807 &GetSEStatus);
809 if (0 != GetSEStatus)
810 {
811 rbe_log_stats(0, "Unable to get overwrite policy.");
812 return NULL;
813 }
815 if (GetSEDestDirTop(SubmitObjID, SubmitElemID,
816 temp_dirtop, SUBMIT_FIELD_MAX,
817 &GetSEStatus) != 0)
818 {
819 rbe_log_stats(0, "Unable to get dirtop.");
820 return NULL;
821 }
823 if (GetSEWorkItemName(SubmitObjID, SubmitElemID,
824 temp_workitem_name,
825 SUBMIT_FIELD_MAX,
826 &GetSEStatus) != 0)
827 {
828 rbe_log_stats(0, "Unable to get work item name.");
829 return NULL;
830 }
832 if (temp_dirtop != NULL)
833 {
834 minus_c = "-C ";
835 minus_c_arg = temp_dirtop;
836 }
838 work_item = rbc_find_workitem_in_config(temp_workitem_name,
```

Page 43 of 144	make_remote_cpiogen_cmd	Thu Jan 03 12:25:21 2008	Page 44 of 144	generate_rcmdpath	Thu Jan 03 12:25:21 2008
839 1 840 1 841 1 842 1 843 1 844 2 845 2 846 2 847 2 848 1 850 1 851 2 852 2 853 2 854 2 856 2 857 2 858 2 860 2 861 2 862 1 864 1 865 2 866 2 869 2 870 2 871 1 873 1 874 1 875 1 876 1 877 1 878 1 879 1 881 1 882 1 883 1 884 1 885 1 886 1 887 1 888 1 890 1 891 1	<pre> NULL, NULL, rcx->rc_config, /* OK */ if (work_item != NULL && DEFAULT_EB_BUFSIZE != work_item->recover_client_bufsize) { sprintf(bufsize_buf, "-A -b -A %d", work_item->recover_client_bufsize); bufsize = bufsize_buf; } switch (temp_overwrite_policy) { case RC_OVERPOL_NO_CLOBBER: noclobber = "-A -onever"; break; case RC_OVERPOL_NEW_CLOBBER: noclobber = "-A -onever"; break; default: break; /* do something better here */ } if (debugmode) { static char debugarg[100]; (void) sprintf(debugarg, "-x /tmp/RBdebug%d", getpid()); remdebugflag = debugarg; } len = strlen(rcmdpath) + strlen(bufsize) + strlen(remdebugflag) + strlen(minus_c) + strlen(minus_c_arg) + strlen(noclobber) + 1; /* for '\0' */ cmd = sm_malloc(len); (void) sprintf(cmd, "%s%s%s%s", rcmdpath, bufsize, remdebugflag, minus_c, minus_c_arg, noclobber); return cmd; } /* end of make_remote_cpiogen_cmd() */</pre>	894 895 896 897 898 899 900 901 902 903 904 906 907 908 909 1 910 1 911 1 912 1 913 1 914 1 915 1 917 1 918 1 920 1 922 1 923 1 924 1 925 2 926 2 927 2 928 1 930 1 931 1 932 1 933 2 934 2 935 2 936 1 937 1 938 1 939 1 940 1 941 1 942 1 943 1 945 1 946 2 947 2 948 3 949 3 950 2 951 1 953 1 954 1 955 1	<pre>/* * 1) Remove rcx references. * 2) Research whether ebc_normalize updates. */ /* * Construct the path name of the remote command that * will be executed on the destination client. * * Return ptr to constructed string. * NOTE: caller must copy if string is to be preserved. */ static char * generate_rcmdpath(int SubmitObjID, int SubmitElemID) { static char *mybuf = NULL; size_t len_needed; char *pph; /* pointer to %h */ char *p; char *q; char *norm_host; char temp_scriptname[SUBMIT_FIELD_MAX]; char temp_client_hostname[SUBMIT_FIELD_MAX]; int GetSEStatus = 0; if (GetSErcmdScriptName(SubmitObjID, SubmitElemID, temp_scriptname, SUBMIT_FIELD_MAX, &GetSEStatus) != 0) { rbe_log_stats(0, "Unable to get rcmd script name."); return NULL; } if (GetSEDestClientName(SubmitObjID, SubmitElemID, temp_client_hostname, SUBMIT_FIELD_MAX, &GetSEStatus) != 0) { rbe_log_stats(0, "Unable to get client destination name."); return NULL; } /* * If there is a %h in the rc_client_scriptname, * that means put the hostname in there. * * Sorry, no escapes implemented. You can't have an * rc_client_scriptname with a literal %h in it. Tough. */ for (pph = temp_scriptname; *pph != '\0'; pph++) { if (*pph == '%' && *(pph+1) == 'h') { break; } } /* * no %h, just use the bare client scriptname */ if (*pph == '\0')</pre>		
Page 43 of 144	RSlauxmgr.c 15	Thu Jan 03 12:25:21 2008	Page 44 of 144	RSlauxmgr.c 16	Thu Jan 03 12:25:21 2008

```
958 2 {
959 2     return temp_scriptname;
960 1 }
962 1 /*
963 1  * there is a %h ... insert the destination client
964 1  * name into the rc_client_scriptname. First
965 1  * compute how much storage will be needed to
966 1  * hold the result.
967 1 */
969 1 if (!ebc_can_it_be_normalized(temp_client_hostname))
970 2 {
971 2     rec_api_log_csm(SUB_CSM_NOMEM, NULL);
972 2     rbe_log_stats(
973 2         0, "Could not allocate memory generate_rcmdpath");
974 2     /* StartDoneCallBack(EP_RB_RECOVER_FATALERR); */
975 1     return NULL;
977 1 }
978 1 norm_host = ebc_normalize(temp_client_hostname);
979 1 len_needed =
980 1     strlen(temp_scriptname) - 2 + /* -2: %h */
981 1     strlen(norm_host) + 1; /* +1: '\0' */
982 1 if (mybuf == NULL || strlen(mybuf) < (size_t)(len_needed-1))
983 2 {
984 2     if (mybuf != NULL)
985 3         free(mybuf);
986 3     }
987 2 }
989 2 if ((mybuf = malloc((int)len_needed)) == NULL)
990 3 {
991 3     rec_api_log_csm(SUB_CSM_NOMEM, NULL);
992 3     rbe_log_stats(
993 3         0, "Could not allocate memory generate_rcmdpath");
994 2     return NULL;
995 1 }
997 1 q = mybuf;
998 1 for (q = mybuf, p = temp_scriptname; p < pph; q++, p++)
999 2 {
1000 2     *q = *p;
1001 1 }
1003 1 (void)strcpy(q, norm_host);
1004 1 (void)strcat(mybuf, pph+2);
1006 1 return mybuf;
1007 1 } /* end of generate_rcmdpath() */
```

```
1009 int
1010 StartWorkItemRestore(struct recover_context *rcx,
1011     struct auxproc *xp,
1012     int SubmitObjectID,
1013     int SubmitElemID)
1014 1 {
1015 1     char *rcmdv[4]; /* 0 hostname, 1 locuser, 2 remuser, 3 cmd */
1016 1     int changed_priv = 0;
1017 1     int xcplog_pid;
1018 1     rcmd_pkt0_info pkt0p_buffer;
1019 1     char *err_str_buffer = NULL;
1021 1     char temp_client_hostname[SUBMIT_FIELD_MAX];
1022 1     char temp_client_rbumame[SUBMIT_FIELD_MAX];
1023 1     long temp_effective_uid;
1024 1     boolean_t temp_src_sysadmin;
1026 1     int GetSEStatus = 0;
1027 1     int GetSOSStatus = 0;
1030 1     if ( GetSEdestClientName(SubmitObjectID, SubmitElemID,
1031 1         temp_client_hostname,
1032 1         SUBMIT_FIELD_MAX,
1033 1         &GetSEStatus) != 0)
1034 2     {
1035 2         rbe_log_stats(0, "Unable to get client destination name.");
1036 2         return -1;
1037 1     }
1041 1     if ( GetSRCClientUserName(SubmitObjectID, SubmitElemID,
1042 1         temp_client_rbumame,
1043 1         SUBMIT_FIELD_MAX,
1044 1         &GetSEStatus) != 0)
1045 2     {
1046 2         rbe_log_stats(0, "Unable to get client user name.");
1047 2         return -1;
1048 1     }
1052 1     temp_src_sysadmin = GetSOSourcesSystemAdmin(SubmitObjectID,
1053 1         &GetSOSStatus);
1054 1     if (0 != GetSOSStatus)
1055 2     {
1056 2         rbe_log_stats(
1057 2             0, "Unable to get the source system admin privileges.");
1058 1         return -1;
1060 1     if ( GetSOEffectiveUID(SubmitObjectID,
1061 1         &temp_effective_uid,
1062 1         &GetSOSStatus) != 0)
1063 2     {
1064 2         rbe_log_stats(0, "Unable to get effective uid.");
1065 2         return -1;
1066 1     }
1069 1     if (0 != (temp_src_sysadmin) || 0 != temp_effective_uid)
1070 2     {
1071 2         set_recovery_privileges(rcx, &changed_priv);
1072 1     }
```

```

1074 1      rcmdv[0] = temp_client_hostname;
1075 1      rcmdv[1] = temp_client_rbuname;
1076 1      rcmdv[2] = temp_client_rbuname;
1077 1      rcmdv[3] = make_remote_cpilogen_cmd(rcx,
1078 1          SubmitObjectID,
1079 1          SubmitElemID);
1081 1      if(NULL == rcmdv[3])
1082 2      {
1083 2          return -1;
1084 2      }
1085 1      xcpiogen_pid = start_cpilogen(rcx, xp,
1086 1          SubmitObjectID,
1087 1          SubmitElemID,
1088 1          0,
1089 1          rcmdv,
1090 1          &pktOp_buffer,
1091 1          &err_str_buffer);

```

```

1093 1      if (changed_priv)
1094 2      {
1095 2          reset_recovery_privileges(rcx, changed_priv);
1096 1      }
1097 1      return xcpiogen_pid;
1098 1      } /* StartWorkItemRestore() */

```

```

1100 1      /*
1101 1      * Check to see if the user has root access to the
1102 1      * destination client. If so, give him/her the root
1103 1      * privileges on the recovery.
1104 1      */
1106 1      static void
1107 1      set_recovery_privileges(struct recover_context *rcx,
1108 1          int *changed)
1109 1      {
1110 1          int root_access;
1111 1          eperno errnum;
1114 1          *changed = 0;
1115 1          (void)rbc_canirecover(rcx->rc_config, rcx->rc_client_hostname,
1116 1              rcx->rc_human_uidname, &root_access,
1117 1              &errnum);
1118 1          if (root_access)
1119 2          {
1120 2              if (debugmode)
1121 3              {
1122 3                  rbe_log_stats(
1123 3                      0, "the user is a sys admin for the dest client");
1124 2              }
1125 2              rcx->rc_recovery_flags |= RC_RECFLAG_DEST_SYSADMIN;
1126 2              if (rcx->rc_effective_uid != 0)
1127 3              {
1128 3                  if (debugmode)
1129 4                  {
1130 4                      rbe_log_stats(0, "changing the uid to admin status");
1131 3                  }
1132 3                  rcx->rc_effective_uid = 0;
1133 3                  rcx->rc_effective_uidname = "root";
1134 3                  *changed = SET_ROOT;
1135 2              }
1136 2              else
1137 3              {
1138 3                  if (debugmode)
1139 4                  {
1140 4                      rbe_log_stats(
1141 4                          0, "effect uid = %d", rcx->rc_effective_uid);
1142 3                  }
1143 3              }
1144 3          }
1145 2          /*
1146 2          * The user does not have root access to the dest
1147 2          * client. But s/he is the system admin for the
1148 2          * source client, therefore, we need to strip the
1149 2          * root stuff from the user during the recovery.
1150 2          */
1151 2          rcx->rc_recovery_flags ^= RC_RECFLAG_DEST_SYSADMIN;
1152 2          if (rcx->rc_recovery_flags & RC_RECFLAG_SOURCE_SYSADMIN)
1153 3          {
1154 3              if (debugmode)
1155 4              {
1156 4                  rbe_log_stats(
1157 4                      0, "changing the uid to regular user status");
1158 3              }

```

```
1160 3 rcx->rc_effective_uid = rcx->rc_human_uid;
1161 3 rcx->rc_effective_uidname = rcx->rc_human_uidname;
1162 3 *changed = SET_USER;
1163 2
1164 1 }
1165 } /* end of set_recovery_privileges() */
```

```
1166 /*
1167  * Reset the user's identity. The user may have root access
1168  * on the destination client, but not on the source, and vice
1169  * versa.
1170  */
1171
1172 static void
1173 reset_recovery_privileges(struct recover_context *rcx,
1174                          int changed)
1175 {
1176     if (changed == SET_ROOT)
1177     {
1178         if (debugmode)
1179         {
1180             rbe_log_stats(
1181                 0, "resetting the uid to regular user status");
1182         }
1183     }
1184
1185     rcx->rc_effective_uid = rcx->rc_human_uid;
1186     rcx->rc_effective_uidname = rcx->rc_human_uidname;
1187 }
1188 else
1189 {
1190     if (debugmode)
1191     {
1192         rbe_log_stats(0, "resetting the uid to sys admin status");
1193     }
1194
1195     rcx->rc_effective_uid = 0;
1196     rcx->rc_effective_uidname = "root";
1197 }
1198 } /* end of reset_recovery_privileges() */
```

```
1202  /**
1203  **
1204  ** FUNCTION DESCRIPTION:
1205  **
1206  **      This function will start the workitem restore termination.
1207  **
1208  **
1209  ** INPUTS:
1210  **      int auxproc_pid -- auxproc's pid.
1211  **
1212  ** RETURN VALUE:
1213  **
1214  **      none
1215  **
1216  **
1217  ** SIDE EFFECTS:
1218  **
1219  **      auxproc sent the USR1 signal, auxproc will send xcpioen the
1220  **      TERM signal to quit the restore. Auxproc will wait until the
1221  **      restore terminates by waiting for the remote command's exit
1222  **      status.
1223  **
1224  **++
1225  **/
1226
1227 void
1228 QuitWorkItemRestore(int auxproc_pid)
1229 {
1230
1231     /* Auxproc will now alert xcpioen by sending the
1232     * TERM signal.
1233     * This will give xcpioen the ability to
1234     * clean up tmpfiles and clean-up sockets.
1235     * xcpioen will commit suicide as a result of
1236     * this signal.
1237     */
1238
1239     /*
1240     * Alert the auxproc that we want out.
1241     * recxpio etc should also die as a result
1242     * of xx_read_or_die_xx() in recx. The SIGUSR1
1243     * should not kill the auxproc itself, but only
1244     * notify auxproc of the diminishing restore.
1245     */
1246
1247     (void)kill(auxproc_pid, SIGUSR1);
1248
1249     if (debugmode)
1250     {
1251         rbe_log_stats(0, "gs %d quitting restore in process",
1252             AUXPROCNAME,
1253             auxproc_pid);
1254     }
1255
1256     /*
1257     * now that we have indirectly killed the xcpioen and alerted the
1258     * auxproc,
1259     * the signal from the auxproc will be picked up by the next
1260     * routine, auxprocsig_handler, it will notify the user of the
1261     * results and does the cleaning up.
1262     */
```

```
1263  */
1264  }
1265  return;
1266  /* QuitWorkItemRestore() */
```

```

1267  /** GetAuxprocResults()
1268  **
1269  ** FUNCTION DESCRIPTION:
1270  **
1271  ** Inherited from auxprocsig_handler().
1272  **
1273  ** This routine is called when there is information appears
1274  ** in the aux-process. The aux-process starts xcpioen and
1275  ** is responsible for "listening" to status coming from xcpioen.
1276  ** When status comes back from xcpioen, the aux-process signals
1277  ** this process, which is trapped by the caller. And finally,
1278  ** this routine is called.
1279  **
1280  ** Args:
1281  ** (I) resfd int -- results file descriptor from auxproc.
1282  ** (O) results -- work item results.
1283  **
1284  ** RETURN VALUE:
1285  ** The number of local and remote status collected from fd.
1286  **
1287  ** SIDE EFFECTS:
1288  ** none
1289  **
1290  **
1291  **
1292  **/

1295  int
1296  GetAuxprocResults(int resfd,
1297                    wi_restore_results *results)
1298  {
1299      int      exit_stat;
1300      char     *resbuf;
1301      int      n_read;
1302      char     c = 0;
1303      int      n_status_read = 0;

1305      while ((n_status_read < 2) && (1 == fd_avail_test(resfd)))
1306      {
1307          resbuf = (char *) &exit_stat;

1309          if (1 != pread_or_warn(resfd, &c, 1, auxproc_comm_warning))
1310          {
1311              return 0;
1312          }

1314          if (debugmode)
1315          {
1316              rbe_log_stats(0, "GetAuxprocResults() called: '%c', c);
1317          }

1319          if (c == 'R')
1320          {
1321              if (1 == fd_avail_test(resfd))
1322              {
1323                  /*
1324                   * the 'R' command is the for the remote process status
1325                   */
1326                  n_read = auxres2(resfd, 'R', sizeof(int), &resbuf);
1327                  if (-1 != n_read)
1328                  {
1329                      results -> remote_exit_status = exit_stat;
1330                      results -> remote_exit_set = TRUE;
1331                      n_status_read++;

```

```

1332      5      if (debugmode)
1333      6      {
1334      6          rbe_log_stats(
1335      5              0, "remote exit status obtained: %d", exit_stat);
1336      4          }
1337      4          else
1338      5          {
1339      5              rbe_log_stats(
1340      5                  0, "Internal error: remote exit status Incomplete.");
1341      4          }
1342      3          }
1343      2          else if (c == 'r')
1344      2          {
1345      3              if (1 == fd_avail_test(resfd))
1346      3              {
1347      4                  /*
1348      4                  * the 'r' command is the for the local process status
1349      4                  */
1350      4                  n_read = auxres2(resfd, 'r', sizeof(int), &resbuf);
1351      4                  if (-1 != n_read)
1352      4                  {
1353      4                      results -> local_exit_status = exit_stat;
1354      5                      results -> local_exit_set = TRUE;
1355      5                      n_status_read++;
1356      5                      if (debugmode)
1357      5                      {
1358      5                          rbe_log_stats(
1359      6                              0, "local exit status obtained: %d", exit_stat);
1360      6                      }
1361      5                      }
1362      4                      else
1363      4                      {
1364      5                          rbe_log_stats(
1365      5                              0, "Internal error: local exit status Incomplete.");
1366      5                      }
1367      4                      }
1368      3                      }
1369      2                      }
1370      2                      /* sleep (1); */
1371      1                      } while(1) */
1372      1                      return n_status_read;
1373      1                      /*return n_status_read; */
1374      1                      /* GetAuxprocResults() */
1375      1                      }

```



```

1379  /*
1380  * KillWorkItemRestore()
1381  *
1382  * Kill the work item restore. Keep in mind this
1383  * may only be done if the work item is not running
1384  * a restore.
1385  *
1386  * This routine also does the waitpid for auxproc.
1387  * The waitpid (ChildDone()) eliminates the defunct auxproc
1388  * process.
1389  *
1390  * If the work item is running then one must do the
1391  * following.
1392  *
1393  * 1) Call quitWorkItemRestore()
1394  * 2) Wait for and read results from the cmd_from pipe.
1395  * 3) Call KillWorkItemRestore()
1396  *
1397  * Args:
1398  *   ap_pid -- auxproc pid.
1399  *   cmd_to -- auxproc cmd pipe.
1400  *
1401  * Returns:
1402  *   int -- Zero for success.
1403  */
1404  int
1405  KillWorkItemRestore(int ap_pid, int cmd_to)
1406  {
1407     int killRet;
1408     int apResult;
1409     char *apName=AUXPROCNAME;
1410     char *databuf = NULL;
1411     int ChildDoneRet;
1412
1413     killRet = kill(ap_pid, SIGTERM);
1414
1415     if (-1 == killRet)
1416     {
1417         rbe_log_stats(
1418             0, "Can't send sigterm to \"%s\": Pid: %d, error = %s\n",
1419             apName, ap_pid, strerror(errno));
1420     }
1421
1422     do
1423     {
1424         ChildDoneRet = ChildDone(ap_pid, &apResult);
1425         if(0 == ChildDoneRet) sleep (1);
1426     }
1427     while(0 == ChildDoneRet);
1428
1429     switch (ChildDoneRet)
1430     {
1431     /* -1 internal error. errno is set.
1432     * 0 child still running.
1433     * 1 child exited.
1434     * 2 child signalled (no core)
1435     * 3 child signalled core file generated.
1436     * 4 child stopped.
1437     */
1438     case(-1):
1439         return -1;
1440     /* no break necessary */
1441     }

```

```

1443     case(1):
1444         rbe_log_stats(0,
1445             "Sigterm did not bring down \"%s\": Pid: %d,
1446             \" it instead exited with = %d\n",
1447             apName, ap_pid, apResult);
1448         break;
1449
1450     case(2):
1451     case(3):
1452         if (SIGTERM != apResult)
1453         {
1454             rbe_log_stats(0,
1455                 "Sigterm did not bring down \"%s\": Pid: %d,
1456                 \" instead killed by signal = %s\n",
1457                 apName, ap_pid,
1458                 strsignal(apResult));
1459             break;
1460         }
1461
1462     case(4):
1463         rbe_log_stats(0,
1464             "Sigterm did not bring down \"%s\": Pid: %d,
1465             \" instead stopped by signal = %s\n",
1466             apName, ap_pid,
1467             strsignal(apResult));
1468         break;
1469
1470     }
1471
1472     /* auxcmdpacket(cmd_to, 'q', 0, databuf); */
1473     return 0;
1474 } /* KillWorkItemRestore() */

```

```
1  /*
2      *
3      *
4      * Copyright (c) 1998, 1999 EMC Corporation. All rights reserved.
5
6      *
7      * Table of Contents:
8      *
9      * static char *rb_getmethod(host, int *concurrency, int *client_type)
10     * int loopw(int fd, char *buf, int nbytes, int (*func)());
11     * int fd_avail_1_wait_incr(int fd);
12     */
13
14     #define _POSIX_SOURCE 1
15     #define E_GRANDFATHER 1
16     #define TIME_STRUCT
17     #include <eb/eb_port.h>
18
19     #include <netdb.h>
20     #include <pwd.h>
21     #include <ctype.h>
22     #include <sys/wait.h>
23     #include <util/esl_select.h>
24     #include <errno/esl_strerror.h>
25     #include <util/esl_limit.h>
26     #include <stdlib.h>
27
28     /* SFP header */
29     #include <cdl/cdl.h>
30     #include <cdl/cdl_server.h>
31
32     #include <ebconfig/rbconfig.h>
33     #include <efifb/ffdb.h>
34     #include <efifb/rbfinclnt.h>
35     #include <ebutil/eb_normalize.h>
36     #include <ebutil/ebsock_if.h>
37     #include <ebutil/ebutil.h>
38     #include <eb/rb_log.h>
39     #include <restore/RBprogmsg.h>
40
41     #include <RSLinterns.h>
42     #include <RSLspxits.h>
43     #include <RSLremfd.h>
44     #include "RSLauxSupp.h"
45
46     /* EDMLINK API */
47     #include <edmlink/edmlink_api.h>
48
49     /* lint ugliness */
50     #ifndef Printf
51     #define Printf (void)printf
52     #endif
53     #endif
54     #define MAX_FD 1024
55     int debugmode = 0;
56     int is_symmpath = 0;
57     int sp_clexitdone = 0;
58
59     static int xcpiogen_pid = -1;
60     static int xcpiogen_prog_fd = -1;
61     int logging_channel = -1;
62
63     int xcpiogen_pipe[2] = {-1,-1};
64     static char wfiledir[] = "/usr/epoch/etc/";
65
66     static char *rb_getmethod()
```

```
67     register char *host, uint_t *concurrency, uint_t *client_type);
68     static void sigusr1_handler(int sigval); /* "attention" signal */
69     static void sigterm_handler(int sigval); /* "sigterm signal" */
70     static int write_CDL_no_eintr(int fd, char *buf, int nbytes);
71     static int read_CDL_no_eintr(int fd, char *buf, int nbytes);
72
73     static int ForwardXcpiogenProgress(int xcpiogen_prog_fd,
74                                         int xcpiogen_prog_fd,
75                                         int restore_engine_prog_fd,
76                                         boolean_t *zero_byte_read);
77
78     static int DemuxXchildren(int progress_fd,
79                               int remote_fd,
80                               char *remote_progname,
81                               int xcpiogen_fd,
82                               int xcpiogen_pid,
83                               int *remote_exit);
84
85     extern int fd_avail_test(int); /* Found in RSLauxSupp.c */
86
87     struct auxproc_context
88     {
89         int ap_my_auxnum;
90         int ap_r_fd;
91         int ap_w_fd;
92         int ap_r_bulk_fd;
93         int ap_w_prog_fd;
94         char ap_cmd;
95         int ap_dataen;
96         char *ap_data;
97         int ap_resultlen;
98         char *ap_resultdata;
99         ushort_t ap_shelltcp_port;
100        int ap_have_shelltcp_port;
101        int ap_socketport;
102        char *ap_sockethost;
103        char *ap_workitem;
104        struct rbc_configs *ap_config;
105        char ap_error_message[4096];
106    };
107
108    /* Read remote stderr output from fd.
109    * Parse it; act on it according to protocol.
110    * Remote exit status is the last thing that comes
111    * back via the remote stderr stream, return the remote
112    * exit status via *exitp.
113    */
114
115    enum input_states
116    {
117        INSTATE_NG,
118        INSTATE_SEARCH_PREFIX0,
119        INSTATE_SEARCH_PREFIXN,
120        INSTATE_GATHER_COOKIE,
121        INSTATE_SEARCH_SUFFIXN,
122        INSTATE_GATHER_STATUS,
123        INSTATE_COPY_TO_STDOUT,
124        INSTATE_NEWLINE
125    };
126
127    static boolean_t parse_remote_stderr_info2(int prog_fd,
128                                                int remote_fd,
129                                                int *exitp,
130                                                char *remhostname,
131                                                boolean_t first_call,
132                                                enum input_states
133                                                *state_ptr,
```

```
130          enum input_states
131              *next_state_ptr,
132              boolean_ty
133              *skipping_leading_whitespace,
134              int *parsePos,
135              int *msgPos);
136
137 static char *recover_size_prefix(struct auxproc_context *cxp);
138 static int *ebr_direct_rcmd(char **ahost, unsigned short import,
139                             struct auxproc_context *cxp,
140                             char *locuser,
141                             char *remuser, char *cmd, int *fd2p);
```

```
141          /*
142           * The auxiliary process(es) communicate with the main
143           * process via a simple protocol run on a pair of pipes.
144           *
145           * The parent writes commands to the auxiliary process.
146           * The format of a command is:
147           *
148           * <cmd><data-len><data>
149           *
150           * where
151           * <cmd> is a one-byte command
152           * <data-len> is an "int",
153           * and indicates the number of <data> bytes
154           * <data> is command-specific data.
155           *
156           * <data-len> may be zero, but must always be present. Therefore, the
157           * minimum command "packet" is five bytes long.
158           *
159           * Result packets are written back to the parent. The packet format is
160           * the same as the command format, though (obviously) the format
161           * of the data in a results packet is usually different from
162           * the format in a command packet.
163           *
164           * Communications are assumed to be error-free. In other
165           * words, pipes are assumed to work correctly.
166           */
167 static int atn_ec; /* count of SIGUSR1s (ATTN signals) */
```

```
171          /*
172           * 'z' prefixes identify top-level routines
173           * called from the auxproc switch.
174           *
175           * There is no significant to the 'z' letter -- I just picked
176           * a letter at random (hah!) to identify the top level funcs.
177           */
178 static void z_rcmdfilter(struct auxproc_context *cxp);
179 static void z_exec_separate_auxproc(struct auxproc_context *cxp,
180                                     char *pathname);
181
```

```

183 1      int main(int argc, char *argv[])
184 1      {
185 1          char *auxproc_initial_delay_str = NULL;
186 1          int auxproc_initial_delay_int = 0;

188 1          if (argc != 9)
189 2          {
190 2              exit(1);
191 1          }

193 1          auxproc_initial_delay_str = getenv("AUXPROC_INITIAL_DELAY");
194 1          if (NULL != auxproc_initial_delay_str)
195 2          {
196 2              auxproc_initial_delay_int = atoi(auxproc_initial_delay_str);
197 2              sleep(auxproc_initial_delay_int);
198 1          }

200 1          debugmode = atoi(argv[6]);

202 1          do_auxproc(atoi(argv[1]), /* auxproc ordinate. */
203 1                      atoi(argv[2]), /* cmd to auxproc pipe */
204 1                      atoi(argv[3]), /* cmd from auxproc pipe */
205 1                      atoi(argv[4]), /* bulk to auxproc pipe */
206 1                      atoi(argv[5]), /* prog from auxproc pipe */
207 1                      argv[0], /* progname */
208 1                      argv[7], /* socket host name */
209 1                      atoi(argv[8])); /* socket port */

211 1          return(0); /* Can we get here ?? */
213 1      } /* End main() */

```

```

215 1          /*
216 1          * Worker-bee loop.
217 1          * Read instructions on r_fd.
218 1          * Write results to w_fd.
219 1          */
220 1          int ncnds_rcvd = 0;

222 1          void
224 1          do_auxproc(int procnum,
225 1                      int r_fd,
226 1                      int w_fd,
227 1                      int r_bulk_fd,
228 1                      int w_prog_fd,
229 1                      char *xname,
230 1                      char *sockethost,
231 1                      int socketport)
232 1          {
233 1              struct auxproc_context ctx;
234 1              struct servant *sp;
235 1              sigset_t empty_set;
236 1              int rotation_size;
237 1              int errnum;
238 1              int index_fd;
239 1              /*
241 1              * close all file descriptors that we do not need
242 1              * we only need the ones passed to us from the restore
243 1              * engine. So we can close anything above stderr and anything
244 1              * that is no equal to what was passed in
245 1              */
246 1              for(index_fd = 3; MAX_FD > index_fd; index_fd++)
247 1              {
248 2                  if ((index_fd == r_fd) ||
249 2                      (index_fd == w_fd) ||
250 2                      (index_fd == r_bulk_fd) ||
251 2                      (index_fd == w_prog_fd))
252 2                  {
253 3                      continue;
254 3                  }
255 3                  (void)close(index_fd);
256 2              }
257 2          }

259 1          memset(&ctx, 0, sizeof(struct auxproc_context));
261 1          /*
262 1          * Prepare sigaction parameters
263 1          */
264 1          sigemptyset(&empty_set);

266 1          /*
268 1          * ignore mild keyboard interrupts (^C); parent handles
269 1          */
270 1          eb_set_signal_handler(SIGINT, SIG_IGN, &empty_set, E_SA_RESTART);

272 1          /*
274 1          * SIGUSR1 used to get our attention when the parent
275 1          * process really wants us to stop what we are doing.
276 1          */
277 1          eb_set_signal_handler(

```

Page 63 of 144	do_auxproc	Thu Jan 03 12:25:21 2008
<pre> 280 1 eb_set_signal_handler(281 SIGUSR1, sigusr1_handler, &empty_set, E_SA_RESTART); 282 1 SIGTERM, sigterm_handler, &empty_set, E_SA_RESTART); 283 1 284 1 * Arguments passed are collected together in 285 1 * this context structure only because that makes 286 1 * it easier to add/change arguments in the future 287 1 * (diddle the structure rather than diddle a zillion 288 1 * calls in the switch statement). 289 1 290 1 ctx.ap_my_auxnum = procnum; 291 1 ctx.ap_r_fd = r_fd; 292 1 ctx.ap_w_fd = w_fd; 293 1 ctx.ap_r_bulk_fd = r_bulk_fd; 294 1 ctx.ap_w_prog_fd = w_prog_fd; 295 1 ctx.ap_sockethost = sockethost; 296 1 ctx.ap_socketport = socketport; 297 1 298 1 /* 299 1 * Will need the shell/tcp service later; if 300 1 * can't find it now we might as well complain now. 301 1 */ 302 1 303 1 ctx.ap_have_shelltcp_port = 0; 304 1 if ((sp = getserverbyname("shell", "tcp")) != NULL) 305 1 { 306 1 ctx.ap_shelltcp_port = (ushort_t)sp->s_port; 307 1 ctx.ap_have_shelltcp_port = 1; 308 1 } 309 1 else 310 1 { 311 1 WriteFmtStringMsg(312 1 ctx.ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0, 313 1 "Error: Cannot find shell/tcp network 314 1 service.\n" 315 1 "Browsing of catalogs may continue, but" 316 1 "the \"start\" command will not work."); 317 1 } 318 1 if (NULL == ctx.ap_config) 319 1 { 320 1 if (NULL == (ctx.ap_config = malloc(sizeof(321 1 struct rbc_configs)))) 322 1 { 323 1 /* We would prefer to log this stuff, but we 324 1 * have not opened logging yet. 325 1 */ 326 1 WriteFmtStringMsg(327 1 ctx.ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0, 328 1 "Could not allocate memory in 329 1 do_auxproc"); 330 1 exit(1); 331 1 } 332 1 if (rbc_parse_config(333 1 NULL /*use the default name*/, &ctx.ap_config, 334 1 RBC_PARSE_DO_NOT_PRESERVE 335 1 RBC_PARSE_APPLY) != 0) </pre>		
<pre> 332 3 { 333 3 /* We would prefer to log this stuff, but we 334 3 * have not opened logging yet. 335 3 */ </pre>		<pre> Page 63 of 144 RSLauxmain.c 7 Thu Jan 03 12:25:21 2008 </pre>

Page 64 of 144	do_auxproc	Thu Jan 03 12:25:21 2008
<pre> 337 3 rec_api_log_csm(SUB_CSM_NO_PARSE_CFG, NULL); 338 3 WriteFmtStringMsg(339 3 ctx.ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0, 340 3 "Auxproc -- Cannot parse configuration 341 3 file"); 342 3 } 343 3 if ((errnum = eb_path_init()) != 0) 344 3 { 345 3 exit(1); 346 3 } 347 3 (void)rbe_log_init("auxproc"); 348 3 349 3 /* 350 3 * Determine the "recoveries.log" file rotation size. 351 3 * Use a default unless this was specified in the config file. 352 3 */ 353 3 rotation_size = RBRECOVER_LOGSIZE; 354 3 if ((ctx.ap_config != (struct rbc_configs *)NULL) 355 3 && (ctx.ap_config->srfile.filename != NULL)) 356 3 { 357 3 rotation_size = ctx.ap_config->srfile.rotation_size; 358 3 } 359 3 (void)rbe_log_add(SRATS_LOGGING, eb_recover_logpath, LOGT_FILE, 360 3 rotation_size, &logging_channel); 361 3 362 3 363 3 /* 364 3 * The actual worker-bee loop 365 3 */ 366 3 for (;;) 367 3 { 368 3 char c; 369 3 int datalen; 370 3 #define SMALL_DATALEN 128 371 3 char buf[SMALL_DATALEN]; 372 3 char *data; 373 3 374 3 /* 375 3 * read command byte 376 3 */ 377 3 pread_or_die(r_fd, &c, 1, _exit); 378 3 379 3 /* 380 3 * read data-len 381 3 */ 382 3 pread_or_die(r_fd, &c, 1, _exit); 383 3 384 3 /* 385 3 * read data-len 386 3 */ 387 3 pread_or_die(r_fd, (char *)&datalen, sizeof datalen, _exit); 388 3 389 3 /* 390 3 * We do the read here for simple commands that do not 391 3 * take much data. 392 3 */ 393 3 if (datalen > SMALL_DATALEN datalen == 0) 394 3 { 395 3 data = NULL; 396 3 } 397 3 else 398 3 { 399 3 if (datalen > SMALL_DATALEN datalen == 0) 400 3 { </pre>		<pre> Page 64 of 144 RSLauxmain.c 8 Thu Jan 03 12:25:21 2008 </pre>

```
401 3 {
402 3     pread_or_die(r_fd, buf, datalen, _exit);
403 3     data = buf;
404 2 }
406 2 ++ncmds_rcvcd;
407 2 ctx.ap_cmd = c;
408 2 ctx.ap_datalen = datalen;
409 2 ctx.ap_data = data;
410 2 ctx.ap_resultlen = 0;
411 2 ctx.ap_resultdata = NULL;
413 2 /*
414 2  * NOTE: This switch statement is in alpha-order
415 2  * (case-folded) to make it easier for humans
416 2  * to find entries (at least, it was in alpha-order
417 2  * when I wrote this comment).
418 2 */
420 2 switch (c)
421 3 {
422 3     /*
423 3     * 'P'
424 3     * Ping command. Used for debugging.
425 3     */
427 3     case 'P':
428 3         ctx.ap_resultlen = ctx.ap_datalen;
429 3         ctx.ap_resultdata = ctx.ap_data;
430 3         break;
432 3     /* 'q'
433 3     * Quit command.
434 3     */
436 3     case 'q':
437 3         rbe_close_logs(logging_channel);
438 3         _exit(0);
440 3     /*
441 3     * 'r'
442 3     * rcmd filter. Fire up a process, typically xcpioegen,
443 3     * with output from the process going to an rcmd
444 3     * connection.
445 3     */
447 3     case 'r':
448 3         z_rcmdfilter(&ctx);
449 3         break;
451 3     /* 'X'
452 3     * Exec separate process version of do_auxproc.
453 3     * Used for debugging.
454 3     * No value returned to parent.
455 3     * If exec fails, auxproc dies.
456 3     */
458 3     case 'X':
459 3         z_exec_separate_auxproc(&ctx, Xname);
460 3         _exit(2);
461 3         break;
463 3     default:
464 3         rec_api_log_csm(SUB_CSM_INV_AUXPROC_CMD, NULL);
465 3         rbe_log_stats(
466 3             0, "auxproc -- invalid command in do_auxproc");

```

```
466 3         exit(1);
467 3         break;
468 2     } /* end of switch */
470 2     /*
471 2     * Simple commands set ap_resultlen and ap_resultdata
472 2     * and we push the results to the parent here.
473 2     */
475 2     if (ctx.ap_resultlen >= 0)
476 3     {
477 3         pwrite_or_die(w_fd, &c, 1, _exit);
478 3         pwrite_or_die(w_fd, (char *)&ctx.ap_resultlen,
479 3             sizeof ctx.ap_resultlen, _exit);
480 3         pwrite_or_die(w_fd, ctx.ap_resultdata,
481 3             ctx.ap_resultlen, _exit);
482 2     }
483 2     #if 0
484 2         if ('r' == c)
485 3         {
486 3             rbe_log_stats(0, "Auxproc(%d) local exit is %d.", getpid(),
487 3                 ctx.ap_resultlen);
489 2         }
490 2     #endif
491 1     } /* end of for loop */
492 1     } /* end of do_auxproc() */

```

```

494  /*
495  * Fire up a process with stdin from parent and output to rcmd.
496  * Protocol traffic between "parent" (main process) and us:
497  *
498  * 'r' parent --> auxproc      contains local & remote cmd info
499  * '0' auxproc --> parent      returns "setup" result, see below
500  * 'R' auxproc --> parent      returns remote exit status (1 int)
501  * 'r' auxproc --> parent      returns local exit status (1 int)
502  *
503  * The '0' result packet describes the results of setting things up.
504  * It contains the following information:
505  *
506  * success/failure : int      [ 0 is success,
507  *                             non-zero is a failure code ]
508  * errnum           : int      [ 0 if no applicable errno code ]
509  * pid              : int      [ process ID of local xcpioen proc ]
510  * msglen           : int      [ length, in bytes,
511  *                             of following str ]
512  *
513  * If things were set up successfully, then two integer zero
514  * values, a non-zero pid, and one more zero (msglen) are
515  * sent in the '0' packet.
516  *
517  * The defined failure codes are:
518  * TBD
519  *
520  * The 'R' result packet will not be sent if the '0' result
521  * indicates that an error occurred. The '0' and 'r' result
522  * packets are always sent.
523  *
524  * Parent passes the following in the 'r' packet <data>:
525  * (string) rcmd-hostname
526  * (string) rcmd-locuser
527  * (string) rcmd-remuser
528  * (string) rcmd-cmd
529  * (int) secuid-val
530  * (int) filter-cmd-fd-info (explained below)
531  * (int) future-flags (flags for future hacks:
532  *                   always zero now)
533  * (string) filter-cmd
534  * (int) filter-cmd-argc
535  * (string) filter-cmd-argv0
536  * (string) ...
537  * (string) filter-cmd-argvN
538  * (int) db API socket info
539  * (string) db API socket host name
540  * (string) filespec from workitem
541  *
542  * If filter-cmd-fd-info is -1, then stdout on the filter cmd
543  * is set up to go to the rcmd.
544  * Otherwise, a separate file descriptor (neither stdout nor stderr)
545  * is set up to go to the rcmd, and the argv element indicated
546  * by filter-cmd-fd-info is used as a sprintf template for passing
547  * the file descriptor number to the filter process.
548  *
549  * For example, if filter-cmd-fd-info is 1, then filter-cmd-argv1
550  * should be a sprintf format string, which will be given to
551  * sprintf along with one integer to pass the file descriptor number
552  * to the filter command.
553  *
554  * Items which are (string) are '\0' terminated.
555  *
556  * Parent is responsible for obeying the built-in size limits:

```

```

555  * no more than 100 argv strings for filter-cmd
556  * no more than 10000 total bytes of auxproc data
557  * no more than 100 bytes of filter-cmd-fd-info argv
558  */
559
560  struct rcmd_pkt0_info
561  {
562  int failcode;
563  int errnum;
564  int pid;
565  int msglen;
566  /* variable length char string message follows */
567  };
568
569  /*
570  * Function to build command line prefix for adding environment
571  * for log truncation
572  */
573
574  static char *
575  recover_size_prefix(struct auxproc_context *cxp)
576  {
577  static char lbuf[128];
578
579  if (cxp->ap_config->crfile.rotation_size != NO_ROTATION)
580  {
581  sprintf(lbuf, "EB_MAX_CLIENT_LOG_SIZE=%d;
582  export EB_MAX_CLIENT_LOG_SIZE;",
583  cxp->ap_config->crfile.rotation_size);
584  }
585  else
586  {
587  memset(lbuf, 0, sizeof(lbuf));
588  }
589
590  return lbuf;
591  }
592  /* recover_size_prefix */

```

```

593 static void
594 z_rcmdfilter(struct auxproc_context *cxp)
595 {
596     char *data;
597     char *rcmd_stuff[4];
598     char *human_name;
599     struct rcmd_pkt0_info pkt0;
600     char *pkt0_errstr = "";
601     static int r_resultdata;
602     int filter_cmd_argc;
603     int filter_cmd_fd_info;
604     int filter_cmd_fdout;
605     char *filter_cmd;
606     char *method;
607     int dbsegno=0;
608     int socket_port;
609     int socket_host[100];
610     char *socket_file;
611     char *wifilename[200];
612     FILE *infofd;
613     static char readmode[] = "r";
614     char cbuff[200];
615     char listener_filename[EB_MAXPATHLEN];
616     char *workitem;
617     int rcmd_stder = -1;
618     int *rcmd_fds;
619     int i,n;
620     int c;
621     char *p2;
622     int wait_result;
623     int fdarg[100];
624     char *filter_cmd_argv[100];
625     char databuf[10000];
626     char *filesec;
627     char sslbuff[200];
628     char holdbuff[1024];
629     char *for_command_with_ssl_parameter;
630     char *hold_rbrclient... command;
631     RBC_WORKITEM *pwi;
632     RBC_WORKGROUP *pwg;
633     boolean_t xcpiogen_still_running;
634     e_errno_t errnum;
635
636     /* EDMLINK handle */
637     ELinkHandlePtr_t = NULL;
638     /* EDMLINK objects */
639     ELinkTargetObjPtr_t = NULL;
640     ELinkUseridObjPtr_t = NULL;
641     ELinkCmdObjPtr_t = NULL;
642     /* EDMLINK global options */
643     unsigned long ELinkOptions = 0;
644     /* EDMLINK status */
645     int ELinkStatus = 0;
646     int ELinkStatus_errno = 0;
647     data = exp->ep_data;

```

```

655     if (data == NULL)
656     {
657         /*
658          * command dispatcher did not read data for us,
659          * there was too much. We must read it.
660          */
661         data = databuf;
662         pread_or_die(cxp->ap_r_fd, data, cxp->ap_datalen, _exit);
663     }
664     /*
665      * extract the 4 rcmd strings.
666      * [3] is the rexcpio cmd.
667      */
668     for (i = 0; i < 4; i++)
669     {
670         rcmd_stuff[i] = data;
671         data += strlen(rcmd_stuff[i])+1;
672     }
673     /*
674      * extract the remote uidname to be used
675      */
676     human_name = data;
677     data += strlen(human_name) + 1;
678     /*
679      * extract the filter_cmd-fd-info
680      */
681     memcpy(filter_cmd_fd_info, data, sizeof(int));
682     data += sizeof(int);
683     /*
684      * skip the flags
685      */
686     data += sizeof(int);
687     /*
688      * extract the filter command
689      */
690     filter_cmd = data;
691     data += strlen(filter_cmd)+1;
692     memcpy(filter_cmd_argc, data, sizeof(filter_cmd_argc));
693     data += sizeof(filter_cmd_argc);
694     for (i = 0; i < filter_cmd_argc; i++)
695     {
696         filter_cmd_argv[i] = data;
697         data += strlen(filter_cmd_argv[i])+1;
698     }
699     filter_cmd_argv[filter_cmd_argc] = NULL;
700     /*
701      * extract db API socket info
702      */
703     memcpy((char *)&socket_port, data, sizeof(int));

```


Page 71 of 144	z_rcmdfilter	Thu Jan 03 12:25:21 2008	Page 72 of 144	z_rcmdfilter	Thu Jan 03 12:25:21 2008
720 1	data += sizeof (int);		780 3	/* Cannot restore. Work item not in eb.cfg (No work item for %s.)",	
722 1	/*		781 3	cxp->ap_workitem);	
723 1	* extract db API socket host name		782 3	rec_api_log_csm(SUB_CSM_NO_WRITEM_CFG, csm_err_msg);	
724 1	*/		783 3	rbe_log_stats(0, "Cannot continue without a work item");	
726 1	socket_file = data;		784 3	exit(1);	
727 1	data += strlen(socket_file) + 1;		785 2	}	
729 1	/*		787 2	for (pwi = pwg->pwlist; NULL != pwi; pwi = pwi->next)	
730 1	* extract filespec		788 3	{	
731 1	*/		789 3	if (0 == strcmp(pwi->name, cxp->ap_workitem))	
733 1	filespec = data;		790 4	{	
734 1	data += strlen(filespec) + 1;		791 4	goto gotit;	
736 1	/*		792 3	}	
737 1	* extract workitem		793 2	}	
738 1	*/		794 1	}	
740 1	workitem = data;		796 1	gotit:	
741 1	cxp->ap_workitem = workitem;		797 1	/*	
742 1	data += strlen(workitem) + 1;		798 1	* connect by direct or rsh methods	
744 1	pkt0.failcode = 0;		799 1	*/	
745 1	pkt0.errnum = 0;		801 1	if ((method = rb_getmethod(rcmd_stuff[0], NULL, NULL)) == NULL)	
746 1	pkt0.pid = -1;		802 2	{	
747 1	pkt0.msglen = -1;		803 2	writeFmtStringMsg(
748 1	pkt0_errstr = "";		804 2	cxp->ap_w_prog_fd, EDMREPROGMSG_AUXPROC_WARNING, 0,	
750 1	/*		805 2	"Unable to get connection method to host \"%s\": "	
751 1	* locate work item in config info		806 2	rcmd_stuff[0]);	
752 1	*/		807 2	method = "rsh";	
754 1	if (NULL == cxp->ap_config)		808 1	}	
755 2	{		810 1	/*	
756 2	if (NULL == (cxp->ap_config = malloc(sizeof(811 1	* now check for a work item override of the connection method	
757 3	struct rbc_configs)))		812 1	*/	
758 3	{		814 1	if (CNCTN_RSH != pwi->connection_type) /* if other than rsh */	
759 3	rbe_log_stats(815 2	{	
760 2	0, "Could not allocate memory in z_rcmdfilter!");		816 2	switch (pwi->connection_type)	
761 2	exit(1);		817 3	{	
762 2	if (rbc_parse_config(818 3	case CNCTN_RSH:	
764 3	NULL /*use the default name */ , &cxp->ap_config,		819 3	method = "rsh";	
765 3	RBC_PARSE_DO_NOT_PRESERVE		820 3	break;	
766 3	RBC_PARSE_APPLY) != 0)		822 3	case CNCTN_EDMLINK:	
767 3	{		823 3	method = "edmlink";	
768 2	rec_api_log_csm(SUB_CSM_NO_PARSE_CFG, NULL);		824 3	break;	
769 1	rbe_log_stats(826 3	case CNCTN_DIRECT:	
771 1	0, "auxproc -- Cannot parse configuration file");		827 3	method = "direct";	
772 2	exit(1);		828 3	break;	
773 2	}		830 3	case CNCTN_SOCKET:	
774 3	for (pwg = cxp->ap_config->pgrouplist; ; pwg = pwg->next)		831 3	method = "socket";	
775 3	{		832 3	break;	
777 3	if (pwg == (RBC_WORKGROUP *)NULL)		834 3	case CNCTN_NETWORK:	
778 3	{		835 3	method = "network";	
779 3	char csm_err_msg[256];		836 3	break;	
	rbe_log_stats(838 3	default:	
	0, "\n * No work item name \"%s\" in configuration file",		839 3	method = "???" ;	
	cxp->ap_workitem);		840 3	break;	
	sprintf(csm_err_msg,		841 2	}	
Page 71 of 144	RSlauxmain.c 15	Thu Jan 03 12:25:21 2008	Page 72 of 144	RSlauxmain.c 16	Thu Jan 03 12:25:21 2008

```

843 2         if (debugmode)
844 3         {
845 3             rbe_log_stats(
0, "using connection method \"%s\" due to work item \"%s\" override",
            method, pwi->name);
846 3         }
847 2     }
848 1     }
850 1     if (NULL != pwi && pwi->recovery_init) /* if a recovery init
command specified */
851 2     {
852 2         (void) system(pwi->recovery_init);
853 1     }
/* OSGsw34952 -- Workaround for STPclose() behavioral defect
If pwi->ssl_groupname is not NULL we are going to be using
SP for the restore.
Shut off the SSL (STP) exit handler.
Environment variable is used here so the call to
CDL_noatexit can be avoided, if desired.
860 1
861 1
862 1     */
864 1     if ( NULL == getenv("AUXPROC_DONOT_PERFORM_CDLNOATEXIT") )
865 2     {
866 2         if ( NULL != pwi->ssl_groupname )
867 3         {
869 3             #if 0
870 3                 if ( NULL != getenv("EDM_STP_LOGGING"))
871 4                 {
872 4                     (void) putenv("STP_LVL=DEBUG2");
873 4                     (void) putenv("STP_EVTLOG_SHARE=OFF");
874 3                 }
875 3             #endif
877 3                 (void) CDL_noatexit();
878 3                 is_sympath = 1;
879 2             }
880 1         }
882 1         if (0 == strcmp(method, "direct"))
883 2         {
884 2             /*
885 2              * direct connection method here
886 2              */
888 2             if (debugmode)
889 3             {
890 3                 rbe_log_stats(0, "Invoking cmd with %s, %s, %s\n",
891 3                     rcmd_stuff[0], rcmd_stuff[1],
892 3                     rcmd_stuff[2], rcmd_stuff[3]);
893 2             }
895 2             strcpy(cbuf, recover_size_prefix(exp));
897 2             if (strlen(cbuf) > (size_t) 0)
898 3             {
899 3                 sprintf(holdbuf, "( %s %s )", cbuf, rcmd_stuff[3]);
900 3                 rcmd_stuff[3] = holdbuf;
901 2             }
903 2             rcmd_fd = ebr_direct_rcmd(
&rcmd_stuff[0], exp->ap_shelltcp_port, exp,
rcmd_stuff[1], rcmd_stuff[2],
rcmd_stuff[3],
904 2

```

```

905 2         if (debugmode)
906 2         {
907 2             rbe_log_stats(
908 3                 0, "AUXPROC: rcmd returned fd %d, stderr fd %d\n",
909 3                 (
NULL != rcmd_fd) ? rcmd_fd[0] : -1, rcmd_stderr);
910 2         }
912 2         if (rcmd_fd == NULL)
913 3         {
914 3             pkt0_failcode = 2;
915 3             pkt0_errstr = "cannot set up remote connection";
917 3             goto send_0;
918 2         }
919 1         }
920 1         else if ( (0 == strcmp(method, "rsh")) || (0 == strcmp(
method, "edmlink")) )
921 2         {
922 2             /* variable for dealing with symm path and cross restore */
924 2             char sslName[CDL_HOST_LENGTH];
925 2             char sslGroup[CDL_STG_LENGTH];
926 2             char **stgNames; /* for returned array of STG names */
927 2             int numGroups = 0; /* for returned number of STG names */
928 2             int addIndex = 0;
929 2             int rc = 0;
930 2             boolean_t symmPathOK = FALSE;
931 2             /*
932 2              * rsh connection method here
933 2              */
935 2             /*
936 2              * Set up the rcmd connection (rsh method) to the client.
937 2              */
939 2             if (! exp->ap_have_shelltcp_port)
940 3             {
941 3                 struct servent *sp;
943 3                 /*
944 3                  * Try to get the port number again.
945 3                  */
947 3                 if ((sp = getservbyname("shell", "tcp")) != NULL)
948 4                 {
949 4                     exp->ap_shelltcp_port = (ushort_t) sp->ss_port;
950 4                     exp->ap_have_shelltcp_port = 1;
951 3                 }
952 3                 else
953 4                 {
954 4                     pkt0_failcode = 1;
955 4                     pkt0_errstr = "shell/tcp service not found";
956 4                     goto send_0;
957 3                 }
958 2             }
960 2             /*
961 2              * check for SSL enabled work item and if present add
962 2              * SSL information to command string as an environment
963 2              * variable--but first check for cross restore and make
964 2              * appropriate adjustments
965 2

```

Page 75 of 144	z_rcmdfilter	Thu Jan 03 12:25:21 2008			
967 2 968 3 969 3 970 4 971 4 972 4 973 4 974 5 975 5 976 5 977 5 978 5 979 5 980 5 981 4 982 4 983 5 984 5 985 5 986 5 987 5 988 5 989 5 990 5 991 5 992 5 993 5 994 5 995 5 996 5 997 5 998 5 999 6 1000 6 1001 6 1002 5 1003 5 1004 4 1005 4 1006 4 1007 4 1008 4 1009 4 1010 3 1011 3 1012 4 1013 4 1014 4 1015 4 1016 4 1017 4 1018 4 1019 3	<pre>if (pwi->ssl_groupname != NULL) { if (!ebc_same_host(rcmd_stuff[0], pwi->sysname)) /* if x-recovery */ { /* Grab the STG entries for this client */ rc = CDL_getavailablegroups(rcmd_stuff[0], &numGroups, &stgNames); if ((numGroups==0) (rc < 0)) { char errsStr[128] ; sprintf(errsStr, "Unable to use SymmPath for cross restore to host %s--using network", rcmd_stuff[0]); rbe_user_error(0, errsStr); /* * Default to network by leaving SP flag set to * False. */ } else { symmPathok = TRUE; /* * The destination client is SP enabled. * to see if the same STG group exists. If we don't * find the same one we will use the first, * so set * that as a default */ memset(sslGroup, 0, CDL_STG_LENGTH); stncpy(sslGroup, stgNames[0], CDL_STG_LENGTH-1); /* Initialize the index */ addindex = numGroups; while (--addindex >= 0) if (0 == strcmp(pwi->ssl_groupname, stgNames[addindex])) { stncpy(sslGroup, stgNames[addindex], CDL_STG_LENGTH-1); continue; } sprintf(sslbuf, "EB_SSL_RECOVER=\"%s\""; export EB_SSL_RECOVER ; sslGroup); } } /* not a cross recovery--just normal SP */ symmPathok = TRUE; memset(sslGroup, 0, CDL_STG_LENGTH); stncpy(sslGroup, pwi->ssl_groupname, CDL_STG_LENGTH-1); sprintf(sslbuf, "EB_SSL_RECOVER=\"%s\""; export EB_SSL_RECOVER ; sslGroup); }</pre>	1021 2 1025 2 1026 3 1027 3 1028 3 1029 3 1030 3 1032 3 1033 3 1034 2 1036 2 1038 2 1039 3 1040 3 1041 4 1042 4 1043 4 1044 3 1045 3 1046 4 1047 4 1048 4 1049 3 1051 2 1052 2 1053 3 1054 3 1055 3 1056 2 1058 2 1061 2 1062 2 1063 2 1065 2 1066 2 1068 2 1069 2 1070 3 1071 3 1072 2 1074 2 1075 3 1076 3 1077 3 1078 2 1080 2 1081 2	<pre> } if (debugmode) { sprintf(cxp->ap_error_message, "Calling ELinkShell with %s, %d, %s, %s", rcmd_stuff[0], cxp->ap_shelltcp_port, rcmd_stuff[1], rcmd_stuff[2], rcmd_stuff[3]); rbe_log_stats(0, "%s", cxp->ap_error_message); cxp->ap_error_message[0] = 0; } stncpy(cbuf, recover_size_prefix(cxp)); if (strlen(cbuf) > (size_t)0) { if (symmPathok == TRUE) { sprintf(holdbuf, "%s %s %s", cbuf, sslbuf, rcmd_stuff[3]); rcmd_stuff[3] = holdbuf; } else { sprintf(holdbuf, "(%s %s)", cbuf, rcmd_stuff[3]); rcmd_stuff[3] = holdbuf; } } else if (symmPathok == TRUE) { sprintf(holdbuf, "(%s %s)", sslbuf, rcmd_stuff[3]); rcmd_stuff[3] = holdbuf; } /* if not Symm Path and cbuf is NULL rcmd_stuff[3] stays unchanged */ /* ** EDMLINK API */ /* set the initial edmlink transport methods */ ELinkOptions = ELINK_SHELL_RCMD ELINK_SHELL_REXEC; /* enable the edmlink transport method */ if (0 == strcmp(method, "edmlink")) { ELinkOptions = ELINK_SHELL_EDMLINK; } if (debugmode) { /* turn on edmlink debugging */ ELinkOptions = ELINK_LOGLVL_DEBUG; } /* initialize the edmlink api */ perlinkhandle = ELinkInitAPI(ELinkOptions);</pre>		
Page 75 of 144	RSLauxmain.c 19	Thu Jan 03 12:25:21 2008	Page 76 of 144	RSLauxmain.c 20	Thu Jan 03 12:25:21 2008

Thu Jan 03 12:25:21 2008		z_rcmdfilter	Page 77 of 144
1083 2	if(NULL == pLinkHandle)		
1084 3	{		
1085 3	sprintf(cxp->ap_error_message,		
1086 3	"ERROR: Could not initialize the EDMLINK API.");		
1088 3	WriteFmtStringMsg(cxp-> ap_w_prog_fd,		
1089 3	EDMREPROGMSG_AUXPROC_ERROR, 0,		
1090 3	"%s\n", cxp->ap_error_message);		
1092 3	rbe_log_stats(0, "%s", cxp->ap_error_message);		
1093 3	return;		
1094 2	}		
1096 2	/* get new target host object */		
1097 2	pLinkTargetObj = ELINKNewTargetObj(pLinkHandle,		
1098 2	rcmd_stuff[0]);		
1100 2	if(NULL == pLinkTargetObj)		
1101 3	{		
1102 3	sprintf(cxp->ap_error_message,		
1103 3	"ERROR: Could not create a new EDMLINK target		
	object.");		
1105 3	WriteFmtStringMsg(cxp-> ap_w_prog_fd,		
1106 3	EDMREPROGMSG_AUXPROC_ERROR, 0,		
1107 3	"%s\n", cxp->ap_error_message);		
1109 3	rbe_log_stats(0, "%s", cxp->ap_error_message);		
1110 3	/* clean up and return */		
1111 3	(void) ELINKDoneAPI(pLinkHandle);		
1112 3	return;		
1113 2	}		
1115 2	/* get new user id object */		
1116 2	pLinkUserObj = ELINKNewEbrUserObj(pLinkHandle,		
1117 2	pLinkTargetObj,		
1118 2	rcmd_stuff[1]);		
1120 2	if(NULL == pLinkUserObj)		
1121 3	{		
1122 3	WriteFmtStringMsg(cxp-> ap_w_prog_fd,		
1123 3	EDMREPROGMSG_AUXPROC_ERROR, 0,		
1124 3	"Could not create a new EDMLINK user id		
1125 3	object.\n");		
1126 3	/* clean up and return */		
1127 3	if(NULL != pLinkTargetObj)		
	(void) ELINKDestroyObj(
	pLinkHandle, pLinkTargetObj);		
1128 3	(void) ELINKDoneAPI(pLinkHandle);		
1129 3	return;		
1130 2	}		
1132 2	/* get new command object */		
1133 2	pLinkCmdObj = ELINKNewCmdObj(pLinkHandle,		
1134 2	pLinkTargetObj,		
1135 2	rcmd_stuff[3]);		
1137 2	if(NULL == pLinkCmdObj)		
1138 3	{		
1139 3	sprintf(cxp->ap_error_message,		
1140 3	"ERROR: Could not create a new EDMLINK command		
	object.");		
1142 3	WriteFmtStringMsg(cxp-> ap_w_prog_fd,		
1143 3	EDMREPROGMSG_AUXPROC_ERROR, 0,		
1144 3	"%s\n", cxp->ap_error_message);		
Thu Jan 03 12:25:21 2008		RSlauxmain.c 21	Page 77 of 144

Thu Jan 03 12:25:21 2008		z_rcmdfilter	Page 78 of 144
1146 3	rbe_log_stats(0, "%s", cxp->ap_error_message);		
1147 3	/* clean up and return */		
1148 3	if(NULL != pLinkTargetObj)		
1149 3	(void) ELINKDestroyObj(
	pLinkHandle, pLinkTargetObj);		
1150 3	if(NULL != pLinkUserObj)		
1151 3	(void) ELINKDestroyObj(
	pLinkHandle, pLinkUserObj);		
1152 3	return;		
1153 3	}		
1154 2			
1156 2	rcmd_fd = rcmd_fds; /* set valid pointer */		
1158 2	ELINKStatus = ELINKShell(pLinkHandle,		
1159 2	pLinkTargetObj,		
1160 2	pLinkUserObj,		
1161 2	pLinkCmdObj,		
1162 2	&rcmd_fd[0],		
1163 2	&rcmd_stder);		
1164 2	ELINKStatus_erno = erno;		
1166 2	rcmd_fd[1] = rcmd_fd[0]; /* this one is bi-directional */		
1168 2	if (debugmode)		
1169 3	{		
1170 3	rbe_log_stats(
	0, "AUXPROC: ELINKShell returned %d, erno %d,"		
	" fd %d, stder fd %d\n", ELINKStatus,		
	ELINKStatus_erno, rcmd_fd[0], rcmd_stder);		
1171 3	}		
1172 3			
1173 2			
1175 2	/*		
1176 2	** EDMLINK: clean up and shut down the edmlink api		
1177 2	*/		
1179 2	if(NULL != pLinkTargetObj)		
1180 2	(void) ELINKDestroyObj(pLinkHandle, pLinkTargetObj);		
1182 2	if(NULL != pLinkUserObj)		
1183 2	(void) ELINKDestroyObj(pLinkHandle, pLinkUserObj);		
1185 2	if(NULL != pLinkCmdObj)		
1186 2	(void) ELINKDestroyObj(pLinkHandle, pLinkCmdObj);		
1188 2	(void) ELINKDoneAPI(pLinkHandle);		
1190 2	if (0 != ELINKStatus)		
1191 3	{		
1192 3	sprintf(cxp->ap_error_message,		
1193 3	"cannot set up remote connection when calling "		
1194 3	"ELINKShell with %s, %d, %s, %s, \"%s\",		
1195 3	rcmd_stuff[0], cxp->ap_shelltcp_port,		
	rcmd_stuff[1],		
	rcmd_stuff[2], rcmd_stuff[3],		
	esl_strerror(
	ELINKStatus_erno), ELINKStatus_erno);		
1196 3			
1197 3			
1199 3	pkt0_failcode = 2;		
1200 3	pkt0_errstr = cxp->ap_error_message;		
1202 3	if (debugmode)		
1203 4	{		
1204 4	rbe_log_stats(
Thu Jan 03 12:25:21 2008		RSlauxmain.c 22	Page 78 of 144

```

1205 3      )
1207 3      goto send_0;
1208 2
1210 2      /* add remainder of Symmetric Path handshaking here
1211 2      * if the we made it through the other Symm Path tests
1212 2      */
1213 2      if (symmPathOK == TRUE)
1214 3      {
1215 3          int newfd; /* fd for SSL sopen call */
1217 3          /*
1218 3          * We're going to grab the short name SSL alias for the
1219 3          * long network name.
1220 3          * Useless unless this is a cross-restore.
1221 3          */
1222 3          because rcmd_stuff[0] == pwi->ssl_groupname on a regular
1223 3          * restore)
1224 3          {
1226 3              memset(ssiName, 0, CDL_HOST_LENGTH);
1227 3              strncpy(ssiName, rcmd_stuff[0], CDL_HOST_LENGTH-1);
1228 4              if ((strlen(rcmd_stuff[0]) >= CDL_HOST_LENGTH) &&
1229 4                  (0 >= CDL_getsslhostname(ssiName, rcmd_stuff[0])))
1230 4                  {
1231 4                      char errstr[128] ;
1232 4                      sprintf(
1233 4                          errStr, "Unable to determine the short SymmPath hostname of %s",
1234 4                          rcmd_stuff[0]);
1235 4                      rbe_user_error(0, errStr);
1236 4                      pkt0_failcode = 2;
1237 4                      pkt0_errstr = "Cross-restore is configured to go
1238 4                          through SP, yet destination host is not"
1239 4                          " properly configured in edm.conf";
1240 4                      goto send_0;
1241 4                  }
1242 3          /*
1243 3          * Establish the SymmPath channel
1244 3          */
1245 3          if ((newfd = CDL_client(rcmd_fd[0],
1246 3                          ssiName,
1247 3                          sslGroup)) < 0)
1248 4          {
1249 4              rbe_user_error(0,
1250 4                  "Unable to open an SSL listener connection
1251 4                  -- CDL_client failed");
1252 4              pkt0_failcode = 2;
1253 4              pkt0_errstr = CDL_errstr(newfd);
1254 4              goto send_0;
1255 4          }
1256 3          /* We connected...Use it.. */
1257 3          rcmd_fd[0] = rcmd_fd[1] = newfd;
1258 2      }
1259 2      }
1260 1      }
1261 1      else if (0 == strcmp(method, "netware"))

```

```

1264 2      {
1265 2          /*
1266 2          * netware connection method here
1267 2          */
1268 2          char buf9[EB_MAXPATHLEN];
1269 2          char targetbuf[EB_MAXPATHLEN];
1270 2          char *csa = "?";
1271 2          char *target = NULL;
1272 2
1273 2          filter_cmd_argv[filter_cmd_argc++] = "-/";
1274 2          /* always full path */
1275 2          filter_cmd_argv[filter_cmd_argc] = NULL;
1276 2          /*
1277 2          * Set up the rcmd connection (rsh method) to the client.
1278 2          */
1279 2          if (! exp->ap_have_shelltcp_port)
1280 2          {
1281 2              struct servant *sp;
1282 2              if (! exp->ap_have_shelltcp_port)
1283 3              {
1284 3                  /*
1285 3                  * Try to get the port number again.
1286 3                  */
1287 3                  if ((sp = getservbyname("shell", "tcp")) != NULL)
1288 3                  {
1289 3                      exp->ap_shelltcp_port = (ushort_t)sp->s_port;
1290 3                      exp->ap_have_shelltcp_port = 1;
1291 3                  }
1292 3                  else
1293 3                  {
1294 3                      pkt0_failcode = 1;
1295 3                      pkt0_errstr = "shell/tcp service not found";
1296 3                      goto send_0;
1297 3                  }
1298 3              }
1299 3          }
1300 3          rcmd_fd = rcmd_fds; /* set valid pointer */
1301 2
1302 2          /*
1303 2          * now skip over ../host/bin/startrec"
1304 2          */
1305 2          for (p2 = rcmd_stuff[3]; '\0' != *p2 && ' ' != *p2; p2++)
1306 2          {
1307 2              /* skip ../host/bin/startrec */
1308 2              while (' ' == *p2)
1309 2              {
1310 2                  p2++; /* advance to start of next word */
1311 2              }
1312 2          }
1313 2          /*
1314 2          * if a -c target:/path then take out target:
1315 2          */
1316 2          *targetbuf = '\0';
1317 2          if (strlen(p2) >= 2)
1318 2          {
1319 2              if (('-' == *p2) && ('/' == p2[1]))
1320 2              {
1321 2                  char *p3 = p2+2;
1322 2                  *targetbuf = p3;
1323 2                  *targetbuf = '\0';
1324 2                  if (strlen(p2) >= 2)
1325 2                  {
1326 2                      if (('-' == *p2) && ('/' == p2[1]))
1327 2                      {
1328 2                          char *p3 = p2+2;

```

```

1329 4      char *p4;
1330 4      while (' ' == *p3)
1331 5      {
1332 5          p3++; /* advance to start of next word */
1333 4      }
1334 4      p4 = p3; /* save this as start of the prefix */
1335 4      for (; '\0' != *p3; p3++)
1336 5      {
1337 5          if ((((' ' == *p3) && ((' ' == p3[1])) ||
1338 5              ((' ' == *p3) && ('\ ' == p3[1]))))
1339 6              break;
1340 6      }
1341 5      if (*p3 != 0) /* if found the prefix */
1342 4      {
1343 4          char * tempChar = NULL;
1344 5          strncpy(targetbuf, p4, p3-p4);
1345 5          targetbuf[p3-p4] = '\0';
1346 5          target = targetbuf;
1347 5          if ('\ ' == p3[1])
1348 5              if ('\ ' == p3[1])
1349 5              {
1350 6                  p3++;
1351 6              }
1352 5          tempChar = strchr(p3+1, '\ ');
1353 5          if (NULL != tempChar)
1354 5              tempChar[0] = ' ';
1355 6          strcpy(p4, p3+1);
1356 6      }
1357 5      }
1358 5      /* now, if it's a netware cross recovery, get password,
1359 4      etc. from
1360 3      * destination target -- don't use source target from pw.
1361 2      * pwi points to work item config struct backup was done
1362 2      under
1363 2      */
1364 2      if (0 != strcmp(
1365 2          rcmd_stuff[0], pwi->sysname) /* if x-recovery */
1366 2      {
1367 2          RBC_WORITEM *pwi2 = pwi;
1368 2      }
1369 2      tsa = "?";
1370 3      if (*targetbuf != 0)
1371 3      {
1372 4          target = targetbuf;
1373 4      }
1374 3      for (pwg = cwp->ap_config->pgrouplist; NULL != pwg;
1375 3          pwg = pwg->next)
1376 3      {
1377 4          for (pwi = pwg->pwlist; NULL != pwi; pwi = pwi->next)
1378 4          {
1379 5              if (ebc_same_host(pwi->sysname, rcmd_stuff[0]))
1380 5              {
1381 6                  goto gotit2;
1382 6              }
1383 5          }
1384 4      }
1385 4      gotit2:
1386 4      }
1387 4      }
1388 4      }
1389 4      }
1390 4      }

```

```

1391 3      /*
1392 3      * if we did not find a wi for the target system,
1393 3      * puc user prompting here.
1394 3      */
1395 3      if (NULL == pwi)
1396 3      {
1397 4          pwi = pwi2; /* restore old value if i found */
1398 4      }
1399 3      }
1400 2      }
1401 2      /*
1402 2      * check for an encrypted password
1403 2      */
1404 2      if (NULL != pwi && (
1405 2          pwi->flags & WORKITEM_FLAGS_ENCRYPT_PASSWD))
1406 3      {
1407 3          sprintf(
1408 3              buf9, "r %s -target %s -tsa %s -login %s -epassword %s", p2,
1409 3              (target != NULL)
1410 3              ? target
1411 3              : ((
1412 3                  pwi->nw_cint_target != NULL) ? pwi->nw_cint_target : "?"),
1413 3              tsa,
1414 3              (pwi->nw_username != NULL) ? pwi->nw_username : "?"),
1415 3              (pwi->nw_passwd != NULL) ? pwi->nw_passwd : "?"));
1416 3          else
1417 3          {
1418 4              sprintf(
1419 4                  buf9, "r %s -target %s -tsa %s -login %s -password %s", p2,
1420 4                  (target != NULL)
1421 4                  ? target
1422 4                  : ((
1423 4                      pwi->nw_cint_target != NULL) ? pwi->nw_cint_target : "?"),
1424 4                  tsa,
1425 4                  (pwi->nw_username != NULL) ? pwi->nw_username : "?"),
1426 4                  (pwi->nw_passwd != NULL) ? pwi->nw_passwd : "?"));
1427 3          }
1428 3      }
1429 3      if (CNCTN_NETWORK != pwi->connection_type)
1430 3      {
1431 4          writeFmtStringMsg(cwp->ap_w_prog_fd,
1432 4              EDMRPROMSG_AUXPROC_WARNING, 0,
1433 4              "work item \"%s\" specifies
1434 4              connection method %s but client
1435 4              "was installed to use the netware
1436 4              method -- using netware method",
1437 4              pwi->name,
1438 4              CNCTN_RSH == pwi->connection_type ? "Rsh" :
1439 4              (
1440 4                  CNCTN_EDMLINK == pwi->connection_type ? "Edmlink" :
1441 4                  (
1442 4                      CNCTN_DIRECT == pwi->connection_type ? "Direct" :
1443 4                      (
1444 4                          CNCTN_SOCKET == pwi->connection_type ? "Socket" :
1445 4                          (
1446 4                              CNCTN_NETWORK == pwi->connection_type ? "Netware" :
1447 4                              "???"
1448 4                          )))
1449 4                      )
1450 4                  )
1451 4              );

```

```
1441 3      }
1443 3      if (0 != pwi->connection_port) /* if port specified */
1444 4      {
1445 4          if (debugmode)
1446 5              rbe_log_stats(
1447 5                  0, "using socket port %d to connect to"
1448 5                  " host \"%s\", witem \"%s\"",
1449 5                  pwi->connection_port, pwi->name);
1450 4      }
1451 4      cxp->ap_shelltcp_port = (ushort_t)pwi->connection_port;
1452 3      }
1453 2      }
1455 2      if (debugmode)
1456 3      {
1457 3          rbe_log_stats(
1458 3              0, "Calling nwr cmd with %s, %d, %s, %s, %s\n",
1459 3              rcmd_stuff[0], cxp->ap_shelltcp_port,
1460 2              rcmd_stuff[1], rcmd_stuff[2], buf9);
1462 2      rcmd_fd[0] = nwr cmd(&rcmd_stuff[0], cxp->ap_shelltcp_port,
1463 2              rcmd_stuff[1], rcmd_stuff[2], buf9,
1464 2              &rcmd_stderr,
1465 2              pwi->ssl_groupname, pwi->ssl_clientname,
1466 2              TRUE);
1467 2      rcmd_fd[1] = rcmd_fd[0]; /* this one is bi-directional */
1468 2      if (debugmode)
1469 3      {
1470 3          rbe_log_stats(
1471 2              0, "AUXPROC: nwr cmd returned fd %d, stderr fd %d\n",
1472 2              rcmd_fd[0], rcmd_stderr);
1473 2      }
1474 2      if (rcmd_fd[0] == -1)
1475 3      {
1476 3          pkt0_failcode = 2;
1477 3          pkt0_errstr = "cannot set up remote connection";
1478 3          goto send_0;
1479 2      }
1480 1      }
1481 1      else if (0 == strcmp(method, "socket"))
1482 2      {
1483 2          /*
1484 2              * socket connection method here
1485 2          */
1487 2      if (CNCTN_RSH != pwi->connection_type) /* if other than rsh */
1488 3      {
1489 3          if (CNCTN_SOCKET != pwi->connection_type)
1490 4          {
1491 4              WriteFmStringMsg(cxp->ap_w_prog_fd,
1492 4                  EDMREPROGMSG_AUXPROC_WARNING, 0,
1493 4                  "work item \"%s\" specifies
1494 4                      connection method %s but "
1495 4                      "client was installed to use the
1496 4                      socket method -- using socket method",
1497 4                      pwi->name,
1498 4                      CNCTN_RSH == pwi->connection_type ? "Rsh" :
1499 4                      "Socket");
```

```
1498 4      }
1499 4      }
1500 4      }
1501 4      }
1502 3      }
1503 3      }
1504 3      }
1505 3      }
1506 3      }
1508 3      socket_port = -1;
1509 3      dbsegno = 1;
1510 3      strcpy(wifilename, wifiledir);
1511 3      if (NULL == socket_file)
1512 4      {
1513 4          rbe_log_stats(
1514 4              0, "No socket file name passed to auxproc\n");
1515 3      }
1516 3      return;
1518 3      }
1519 3      strcat(wifilename, socket_file);
1520 3      /*
1521 3      * Get socket information from eblistend info file
1522 3      */
1523 3      ernum = eb_parse_listener_info_file (wifilename,
1524 3          cbuf,
1525 3          socket_host,
1526 3          &socket_port,
1527 3          &dbsegno,
1528 3          listener_filename);
1529 4      if (ernum != E_SUCCESS)
1530 4      {
1531 4          /*
1532 4              * The routine already logged an error message,
1533 4              * simply
1534 4              * return the error to caller
1535 4          */
1536 4          rbe_log_stats(0,
1537 4              "Unable to read eblistend info file
1538 4                  \"%s\"\n",
1539 4                  wifilename);
1540 3      }
1541 3      return;
1542 4      if (socket_port == -1)
1543 4      {
1544 4          rbe_log_stats(
1545 4              0, "Invalid port field seen for client\n");
1546 3      }
1547 3      if (
1548 4          0 != pwi->connection_port) /* if port specified in work item */
1549 4      {
1550 4          if (debugmode)
1551 5          {
1552 5              rbe_log_stats(
1553 5                  0, "using socket port %d to connect to"
1554 5                  " host \"%s\", witem \"%s\"",
1555 5                  pwi->connection_port, pwi->sysname,
```

```

    }
    }
    socket_port = pwi->connection_port;
}

if (debugmode)
{
    rbe_log_stats(
        ost= %s port=%d seq#=%d\n", socket_host, socket_port, dbseqno);
}

rcmd_fd = rcmd_fds; /* set valid pointer */

/* We need to fix this up...
   of socket (STP) */
Modify sopen to call the right type

rcmd_fd[0] = sopen(socket_host, 'c', socket_port);
if (rcmd_fd[0] < 0)
{
    rbe_log_stats(0, "Socket Open error: %d\n", rcmd_fd[0]);
    _exit(12);
}

}

/*
 * set up ONLY remote command file descriptors to talk to
 * via socket connection. DO NOT overwrite filter_cmd_fds
 */
rcmd_fd[1] = dup(rcmd_fd[0]);

/* If we are connected,
   determine if the user wants normal sockets
   * or ssl.
   */
if (rcmd_fd[0] != -1)
{
    if (!pwi->ssl_groupname)
    {
        /*
         * send OK status to client waiting on data socket
         * data stream
         */
        sprintf(cbuf, "OK: %d socket connection made;
            ", dbseqno);
        (void)write (rcmd_fd[0], cbuf, (int)strlen(cbuf));
    }
    else /* ssl connect */
    {
        char sslName[CDL_HOST_LENGTH];
        char sslGroup[CDL_STG_LENGTH];
        char **stgNames;

        /* for returned array of STG names */
        int numGroups = 0;
        int addIndex = 0;
        int rc = 0;

        /* First test to see if this is a cross restore. If
         * so we can't just use the STG group from the source
         * client
        */
    }
}

```

```

        * since it may not exist on the destination client.
        * We need to see if the destination is SP enabled and
        * get a good STG group for that.

        * if it is available.
            If no STG group is available we will
        * fall back to network.
    */

    if (!ebc_same_host(
        socket_host, pwi->sysname)) /* if x-recovery */
    {
        /* Grab the STG entries for this client */
        rc = CDL_getavailablegroups(
            socket_host, &numGroups, &stgNames);
        if ((numGroups==0) || (rc < 0))
        {
            char errMsg[128] ;
            sprintf(
                errMsg, "Unable to use SymmPath for cross restore to host %s--using
                network", socket_host);
            rde_user_error(0, errMsg);
        }
        /*
        * Default to network and send normal
        * handshake.
        * send OK status to client waiting on data
        * socket before
        */
        * data stream
        */
        sprintf(cbuf, "OK: %d socket connection made;
            ", dbseguo);
        (void)write (rcmd_fd[0], cbuf, (int)strlen(
            cbuf));
    }
    else
    {
        /*
        * The destination client is SP enabled.
        * First check
        * to see if the same STG group exists.
        * If we don't
        * find the same one we will use the first,
        * so set
        * that as a default
        */
        memset(sslGroup, 0, CDL_STG_LENGTH);
        strncpy(
            sslGroup, stgNames[0], CDL_STG_LENGTH-1);
        /* Initialize the index */
        addindex = numGroups;
        while (--addindex >= 0)
            if (0 == strcmp(
                pwi->ssl_groupname, stgNames[addindex]))
            {
                strncpy(
                    sslGroup, stgNames[addindex], CDL_STG_LENGTH-1);
                continue;
            }
        /*
        * We need to get a short name SSL alias if
        the destination
    */

```


Page 87 of 144	z_rcmdfilter	Thu Jan 03 12:25:21 2008
1658 6	<i>/* name happens to be long</i>	
1659 6	<i>*/</i>	
1661 6	memset(sslName, 0, CDL_HOST_LENGTH);	
1662 6	strcpy(sslName, socket_host, CDL_HOST_LENGTH-1);	
1664 6	if ((strlen(socket_host) >= CDL_HOST_LENGTH) && (0 >= CDL_getsslhostname(sslName, socket_host)))	
1667 7	{	
1668 7	char errMsg[128];	
1669 7	sprintf(
1670 7	errMsg, "Unable to determine the short SymmPath hostname of %s",	
1671 7	socket_host);	
1672 7	rbe_user_error(ffdb_errnum, errMsg);	
1673 7	configured to go through SP, yet destination host is not"	
1674 7	" properly configured in edm.conf";	
1675 7	goto send_0;	
1676 6	} if (ebsock_server_connect_ssl(sslGroup,	
1677 6	sslName,	
1678 6	&rcmd_fd[0],	
1679 6	&rcmd_fd[1],	
1680 6	socket_host,	
1681 6	dbsegno,	
1682 6	holdbuff) < 0)	
1683 7	{	
1684 7	writeFmtStringMsg(cxp->ap_w_prog_fd,	
1685 7	EDMREPROGMSG_AUXPROC_ERROR, 0,	
1686 7	<i>/* "%s", */ holdbuff);</i>	
1687 7	return;	
1688 6	}	
1689 5	} else	
1690 4	{	
1691 4	<i>/*</i>	
1692 5	<i>Just a normal SP restore--nothing special to</i>	
1693 5	<i>do</i>	
1694 5	<i>*/</i>	
1695 5	<i>*/</i>	
1697 5	if (ebsock_server_connect_ssl(
1698 5	pwi->ssl_groupname,	
1699 5	pwi->ssl_clientname,	
1700 5	&rcmd_fd[0],	
1701 5	&rcmd_fd[1],	
1702 5	socket_host,	
1703 5	dbsegno,	
1704 6	holdbuff) < 0)	
1705 6	{	
1706 6	writeFmtStringMsg(cxp->ap_w_prog_fd,	
1707 6	EDMREPROGMSG_AUXPROC_ERROR,	
1708 5	0,	
1709 4	<i>/* "%s", */ holdbuff);</i>	
1710 3	return;	
	}	

Page 88 of 144	z_rcmdfilter	Thu Jan 03 12:25:21 2008
1711 2	}	
1713 2	<i>/* Now that we are finally connected, dup to std err */</i>	
1714 2	rcmd_stdcrr = CDL_dup(rcmd_fd[0]);	
1716 1	} else	
1717 1	{	
1718 2	<i>/*</i>	
1719 2	<i>unknown connection method here</i>	
1720 2	<i>*/</i>	
1721 2	writeFmtStringMsg(
1722 2	cxp->ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0,	
1723 2	"Bad connection method \"%s\\n", method);	
1724 2	return;	
1725 1	}	
1727 1	if(0 != pipe(xcplogen_pipe))	
1728 2	{	
1729 2	writeFmtStringMsg(
1730 2	cxp->ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0,	
	"Could not set up progress channel from	
	xcplogen.\\n");	
1731 2	_exit(2);	
1732 1	}	
1734 1	switch (xcplogen_pid = pkt0.pid = fork())	
1735 2	{	
1736 2	case -1:	
1737 2	<i>/* error */</i>	
1738 2	(void) close(xcplogen_pipe[0]);	
1739 2	(void) close(xcplogen_pipe[1]);	
1740 2	pkt0_errstr = "fork failed";	
1741 2	pkt0_failcode = 3;	
1742 2	pkt0_errnum = errno;	
1743 2	break;	
1745 2	case 0:	
1746 2	<i>/* child */</i>	
1748 2	<i>/*</i>	
1749 2	<i>Make our stdin be the bulk-from-parent stream.</i>	
1750 2	<i>*/</i>	
1751 2	(void)dup2(cxp->ap_r_bulk_fd, 0);	
1753 2	if (debugmode)	
1754 3	{	
1755 3	rbe_log_stats(
1756 3	0, "FILTER CMD: \"%s\\n", human_name: \"%s\\n",	
1757 3	filter_cmd, human_name);	
1758 4	for (i = 0; i < filter_cmd_argc; i++)	
1759 4	{	
1760 3	rbe_log_stats(0, "\\t%s\\n", filter_cmd_argv[i]);	
1761 2	}	
1763 2	<i>/*</i>	
1764 2	<i>If no fd info passed to child,</i>	
1765 2	<i>then child stdout goes to the</i>	
1766 2	<i>rcmd process. Otherwise, arbitrary fd is used,</i>	
1767 2	<i>and we pass</i>	
1769 2	<i>the fd number as an argument per the format given to us</i>	
1770 3	<i>*/</i>	
	if (filter_cmd_fd_info == -1)	
	{	

```

1771 3      filter_cmd_fdout = 1;
1772 3      (void)dup2(rcmd_fd[0], 1);
1773 2      }
1774 2      else
1775 3      {
1776 3          filter_cmd_fdout = rcmd_fd[0];
1777 3          (void) sprintf(
1778 3              fdarg, filter_cmd_argv(filter_cmd_fd_info),
1779 3              filter_cmd_fdout);
1780 2          filter_cmd_argv(filter_cmd_fd_info) = fdarg;
1781 2      }
1782 2      /*
1783 2      * Send real uid string on filter_cmd_fdout to client process
1784 2      */
1785 2      i = (int)strlen(human_name);
1786 2      if ((n = looprw(
1787 2          filter_cmd_fdout, human_name, i, write_CDL_no_eintr)) != i)
1788 3      {
1789 3          rbe_log_stats(RBRCOVER_MKERR(
1790 3              errno), "Can't write %d bytes to host \"%s\", "
1791 3              "ec=%d, errno=%d, fd=%d\n", i, rcmd_stuff[0],
1792 3              n, errno,
1793 3              filter_cmd_fdout);
1794 3          _exit(2);
1795 2      }
1796 2      if ((n = looprw(
1797 2          filter_cmd_fdout, "\n", 1, write_CDL_no_eintr)) != 1)
1798 3      {
1799 3          rbe_log_stats(RBRCOVER_MKERR(errno),
1800 3              "Can't write %d bytes to host \"%s\", "
1801 3              "ec=%d, errno=%d, fd=%d\n",
1802 3              1, rcmd_stuff[0], n, errno, filter_cmd_fdout);
1803 3          _exit(3);
1804 2      }
1805 2      (void)close(xcpio_gen_pipe[0]);
1806 2      (void)dup2(xcpio_gen_pipe[1], 2);
1807 2      (void)close(xcpio_gen_pipe[1]);
1808 2      /* Call the CDL layer so we can warn SSL that an execvp
1809 2      is coming. If this isn't an SSL socket this call is a
1810 2      no-op
1811 2      */
1812 2      (void)CDL_execvpPrep(filter_cmd_fdout);
1813 2      (void)execvp(filter_cmd, filter_cmd_argv); /* run xcpio_gen */
1814 2      rbe_log_stats(RBRCOVER_MKERR(
1815 2          errno), "Can't exec \"%s\", errno=%d\n",
1816 2          filter_cmd, errno);
1817 2      _exit(1);
1818 2      break;
1819 2      default:
1820 2      /* Parent */
1821 2      (void)close(xcpio_gen_pipe[1]);
1822 2      xcpio_gen_prog_fd = xcpio_gen_pipe[0];
1823 2      break;
1824 2      } /* end of switch */
1825 2      /* Parent has no use for this file descriptor any more,
1826 2      and nuking it
1827 2      */
1828 2      _exit(1);
1829 1

```

```

1830 1      * now is important so that if the child dies while the rcmd still
1831 1      * wants more input from the child, the rcmd will die too (
1832 1      * hanging around hoping that maybe this parent process will
1833 1      * the data) */
1834 1      /* we didn't really dup these if they were SFP sockets so
1835 1      avoid the closes in that case
1836 1      */
1837 1      if (!pwi->ssl_groupname)
1838 2      {
1839 2          if (rcmd_fd[0] != -1)
1840 2          {
1841 2              (void)close(rcmd_fd[0]);
1842 2          }
1843 2          if (rcmd_fd[1] != -1)
1844 2          {
1845 2              (void)close(rcmd_fd[1]);
1846 2          }
1847 2      }
1848 1      /* However in the case of nwrcmd we did get a network
1849 1      socket for rcmd_stderr so we do need to issue one
1850 1      CDL_close on the one SFP socket. We can detect that
1851 1      case because rcmd_stderr will not be equal to rcmd_fd[0]
1852 1      */
1853 1      else
1854 1      {
1855 1          if (rcmd_stderr != rcmd_fd[0])
1856 2          {
1857 2              if (rcmd_fd[0] != -1)
1858 3              {
1859 3                  /*
1860 3                  * IF not a SSL socket, do the close if not
1861 3                  * we close SP socket in XCPIOGEN after
1862 3                  * data has been moved by calling STPExit()
1863 3                  */
1864 3                  if ( 0 == CDL_issSLsocket(rcmd_fd[0]) )
1865 4                  {
1866 4                      (void) CDL_close(rcmd_fd[0]);
1867 4                  }
1868 4                  rcmd_fd[0] = -1;
1869 4              }
1870 2          }
1871 2          if ( ( 0 == strcmp(method, "rsh") ) || ( 0 == strcmp(
1872 2              method, "edmlink") ) )
1873 2          {
1874 2              /* Because of occasional loss of data, due to timing issues when
1875 2              * connected via edmlink,
1876 2              * also be addressed by edmlink, directly. (OSGSW38596)
1877 2              */
1878 2              int on = 1;
1879 2              if (-1 == setsockopt( rcmd_stderr,
1880 2                  SOL_SOCKET,
1881 2                  SO_KEEPALIVE,
1882 2                  (char *) &on,
1883 2                  sizeof(on) ) )
1884 3              {
1885 3                  rbe_log_stats( RBRCOVER_MKERR(errno),
1886 3                      "Warning: setsockopt for SO_KEEPALIVE
1887 3                      failed" );
1888 3              }
1889 3          }
1890 1

```

Page 91 of 144	z_rcmddfilter	Thu Jan 03 12:25:21 2008
1891 2	}	
1892 1		
1894 1	if(0 == pkt0.failcode)	
1895 2	{	
1896 2	/* Lets check here to see if xcpioegen is still running.	
1897 2	/* The child could have exited prior to the fork.	
1898 2	int ChildDone_ret;	
1899 2	int ChildExitStatus;	
1900 2		
1902 2	r_resultdata = 0; /* "meaningless" */	
1904 2	xcpioegen_still_running = TRUE;	
1905 2	sleep (1);	
1906 2	ChildDone_ret = ChildDone(xcpioegen_pid, &ChildExitStatus);	
1907 2	if(-1 == ChildDone_ret)	
1908 3	{	
1909 3	rbe_log_stats(RBRECOVER_MKERR(errno),	
1910 3	"Internal error: testing to see if xcpioegen	
1911 2	exited.");	
1912 2	} else if (0 != ChildDone_ret)	
1913 3	{	
1914 3	xcpioegen_still_running = FALSE;	
1915 3	r_resultdata = ChildExitStatus;	
1916 2	}	
1917 1	/*	
1918 1	*/	
1919 1	/* Send the setup failure/success packet.	
1920 1	*/	
1922 1	send_0:	
1923 1	if (pkt0.msglen < 0)	
1924 2	{	
1925 2	pkt0.msglen = (int)strlen(pkt0_errstr) + 1;	
1926 1	}	
1928 1	c = '0';	
1929 1	i = (int)sizeof(pkt0) + pkt0.msglen;	
1930 1	write_or_die(cxp->ap_w_fd, &c, 1, _exit);	
1931 1	write_or_die(cxp->ap_w_fd, (char *)i, sizeof i, _exit);	
1932 1	write_or_die(cxp->ap_w_fd, (char *)&pkt0, sizeof pkt0, _exit);	
1934 1	if (pkt0.msglen > 0)	
1935 2	{	
1936 2	sprintf(cxp->ap_error_message, "%s", pkt0_errstr);	
1937 2	rbe_log_stats(0, "%s", cxp->ap_error_message);	
1938 2	write_or_die(cxp->ap_w_fd, pkt0_errstr, pkt0.msglen, _exit);	
1939 1	}	
1941 1	/* If the fork was successful,	
1942 1	wait for the remote exit status to come back	
1943 1	on stderr,	
1944 1	and send it to parent in an 'R' reply. If the fork was	
1945 1	unsuccessful, there is no 'R' reply, and, furthermore,	
1946 1	the value in the	
1947 1	'R' reply is meaningless. */	
1948 1	if ((xcpioegen_pid > 0) &&	
1949 2	(TRUE == xcpioegen_still_running))	
1950 2	{	
1951 2	int remote_exitinfo;	
1952 2	int Demux_ret = 0;	
1953 2	Demux_ret = DemuxAuxChildren(cxp->ap_w_prog_fd,	

Page 92 of 144	z_rcmddfilter	Thu Jan 03 12:25:21 2008
1953 2	rcmd_stder, rcmd_stuff[0],	
1954 2	xcpioegen_prog_fd,	
1955 2	xcpioegen_pid,	
1956 2	&remote_exitinfo);	
1957 2	if(0 != Demux_ret)	
1958 2	{	
1959 3	/* rbe_log_stats(
1960 3	0, "Error monitoring Auxproc's children."); */	
1961 3	/* error_logging */	
1962 2	}	
1965 2	/*	
1966 2	*/	
1967 2	/* notify the main program of remote success/failure	
1969 2	c = 'R';	
1970 2	i = sizeof(remote_exitinfo);	
1971 2	write_or_die(cxp->ap_w_fd, &c, 1, _exit);	
1972 2	write_or_die(cxp->ap_w_fd, (char *)i, sizeof i, _exit);	
1973 2	write_or_die(cxp->ap_w_fd, (
1974 2	char *)&remote_exitinfo, i, _exit);	
1975 2	#if 0	
1976 2	rbe_log_stats(0, "Auxproc(%d) remote exit is %d.", getpid(),	
1977 2	remote_exitinfo);	
1978 2	#endif	
1979 2	/*	
1980 2	*/	
1982 2	/* Is our connection method anything BUT "netware"?	
1983 3	if (0 != strcmp(method, "netware"))	
1984 3	{	
1985 3	/* Now wait for the exit code of the filter (
1986 3	local) process	
1987 3	*/	
1988 3	while ((waitpid(xcpioegen_pid, &wait_result, 0) == -1) &&	
1989 3	(EINTR == errno))	
1991 3	{	
1992 3	/* empty while loop */	
1993 4	}	
1994 4	/*	
1995 3	*/	
1997 3	/* exited while loop either because the return value of	
1998 3	waitpid is NOT -1, or the errno is NOT EINTR	
1999 3	*/	
2000 3	else	
2001 2	/* IT'S NETWARE */	
2002 2	{	
2003 3	/*	
2004 3	***** HACK ***** HACK ***** HACK	
2005 3	***** HACK ***** HACK ***** HACK	
2006 3	***** HACK ***** HACK ***** HACK	
2007 3	***** HACK ***** HACK ***** HACK	
2008 3	***** HACK ***** HACK ***** HACK	
2009 3	***** HACK ***** HACK ***** HACK	
2010 3	***** HACK ***** HACK ***** HACK	

```

2011 3 * This hack is basically for Netware.
2012 3 * does not send anything back to the ( Netware TCP/IP close()
2013 3 * that is running as pid pkt0.pid. local)filter process
2014 3 * just keeps on sending data across the network and
2015 3 * anybody there to read it.
2016 3 *
2017 3 * Therefore, this "auxproc" doesn't finish until all data has
2018 3 * been written to the socket even though no process is
2019 3 * reading
2020 3 * it. However, the process might hang indefinitely if the network
2021 3 * buffer(s) become(s) full.
2022 3 *
2023 3 * This is how we will get around this ( kind of a hack but could
2024 3 * not think of a better way). None of the normal TCP/IP
2025 3 * mechanisms seem to work with Netware.
2026 3 *
2027 3 * If the exitinfo ( which is zero for success or non-zero for
2028 3 * failure) reports failure...
2029 3 * After we have received the exit info from the Netware
2030 3 * then try to get it's "wait_result" twice with a 2
2031 3 * in between. If we don't get it, second pause
2032 3 * process is hung and can't wind down.
2033 3 *
2034 3 * Therefore, let's send a SIGPIPE to the process. I chose
2035 3 * SIGPIPE because this is usually trapped by our
2036 3 * handled in a controlled manner.
2037 3 * it, sleeping for 2 seconds, Just keep on trying to kill
2038 3 * it DOES go away. and then see if it went away until
2039 3 * Just keep on trying to kill it, sleeping for 2 seconds, and
2040 3 * then see if it went away until it DOES go away.
2041 3 *
2042 3 * Apparently, the kill( SIGPIPE) is ignored by the process while
2043 3 * there is activity occurring like positioning the tape.
2044 3 * the reason for having to do the kill( That is
2045 3 * works. ) in a loop until it
2046 3 *
2047 3 *
2048 3 *
2049 3 *
2050 3 if (remote_exitinfo)
2051 3 {
2052 3 /**
2053 3 ** We recieved an error exit status from Novell
2054 3 **/
2055 3 int exitinfo;

```

```

2057 4 exitinfo = waitpid(
2058 4 xcpio_gen_pid, &wait_result, WNOHANG);
2059 4 while (!exitinfo)
2060 4 {
2061 4 /*
2062 4 * Wait a couple of seconds longer and try again
2063 4 */
2064 4 sleep(2);
2065 4 exitinfo = waitpid(
2066 4 xcpio_gen_pid, &wait_result, WNOHANG);
2067 4 if (!exitinfo)
2068 4 {
2069 4 kill(xcpio_gen_pid, SIGPIPE); /* BLAST IT! */
2070 4 }
2071 4 }
2072 4 }
2073 4 else
2074 4 {
2075 4 /**
2076 4 ** We received a success return status from Novell
2077 4 **/
2078 4 *
2079 4 * waitpid() may return with EINTR before the child
2080 4 * proc exits, in that case, we want to continue
2081 4 * with the wait. (fix for OSGsw15487)
2082 4 */
2083 4 while ((waitpid(
2084 4 xcpio_gen_pid, &wait_result, 0) == -1) &&
2085 4 (EINTR == errno))
2086 4 {
2087 4 /* empty while loop */
2088 4 }
2089 4 *
2090 4 * exited while loop either because the return value
2091 4 * of
2092 4 * waitpid is NOT -1, or the errno is NOT EINTR
2093 4 */
2094 4 }
2095 4 }
2096 4 *
2097 4 * main command loop in caller will send this result for us
2098 4 */
2099 4 *
2100 4 *
2101 4 /*
2102 4 * wait_result contains the exit status of the child
2103 4 * process. Use the standard macros to determine the
2104 4 * status of the child process and set the return value
2105 4 * to indicate the status accurately.
2106 4 */
2107 4 if (WIFEXITED(
2108 4 wait_result)) /* child proc exited normally */
2109 4 {
2110 4 /*
2111 4 * set the result to the exit code
2112 4 */
2113 4 wait_result = WEXITSTATUS(wait_result);
2114 4 }
2115 4 }

```

```

2116 2         else if (WIFSIGNALED(wait_result))
2117 3         {
2118 3             /*
2119 3              * child proc exited due to an uncaught signal, so
2120 3              * set the exit code to the signal number plus our own
2121 3              * code XG_EXIT_SIGBASE.
2122 3              */
2123 3             wait_result = WTERMSIG(wait_result) + XG_EXIT_SIGBASE;
2124 3         }
2125 2         else if (WIFSTOPPED(wait_result))
2126 2         {
2127 3             /*
2128 3              * child proc is stopped, set the exit code accordingly.
2129 3              */
2130 3             wait_result = XG_EXIT_STOPPED;
2131 3         }
2132 3
2133 2
2134 2         r_resultdata = wait_result;
2135 2
2136 1         if (rcmd_stderx != -1)
2137 1         {
2138 1             /* OSGsw34952 -- Workaround for STPclose() behavioral defect
2139 2             {
2140 2                 * Shut off the SSL (STP) exit handler
2141 2                 * IF not a SSL socket, do the close if not
2142 2                 * we close SP socket in XCPROGEN after
2143 2                 * data has been moved by calling STPexit()
2144 2                 */
2145 2                 if ( 0 == CDL_isslsocket(rcmd_stderx) )
2146 2                 {
2147 2                     (void)close(rcmd_stderx);
2148 2                     rcmd_stderx = -1;
2149 2                 }
2150 2                 else /* SSL socket used as stderx, so clean up
2151 2                     the SP stuff */
2152 2                 {
2153 2                     /* SP socket was set to stderx of remote
2154 2                     commands. Since, we call CDL_noatexit
2155 2                     when we created the SP socket we will
2156 2                     explicitly call CDL_exit() that will
2157 2                     clean up the entire SP environment */
2158 2                     if ( 0 == sp_clexitdone )
2159 2                     {
2160 2                         CDL_exit();
2161 2                         sp_clexitdone = 1;
2162 2                     }
2163 2                 }
2164 2             }
2165 2
2166 2             /*
2167 2              * set up variables that main loop will use to send 'r' reply
2168 2              */
2169 2             exp->ap_resultlen = sizeof r_resultdata;
2170 2             exp->ap_resultdata = (char *)&r_resultdata;
2171 2             /* end of z_rcmdfilter() */
2172 2         }
2173 2
2174 1
2175 1

```

```

2178         static enum input_states decodecookie(char *cp);
2179
2180         struct cookie2state {
2181             char *cookie;
2182             enum input_states state;
2183         };
2184
2185         static struct cookie2state c2stbl[] = {
2186             { REMFD_COOKIE_FD2OUT,          INSTATE_COPY_TO_SNDOUT },
2187             { REMFD_COOKIE_FD2OUT2,          INSTATE_COPY_TO_SNDOUT },
2188             { REMFD_COOKIE_FD2OUTEND,        INSTATE_SEARCH_PREFIX0 },
2189             { REMFD_COOKIE_STATUS,           INSTATE_GATHER_STATUS },
2190             { NULL,                          INSTATE_NG }
2191         };
2192
2193         static enum input_states
2194         decodecookie(char *cp)
2195         {
2196             struct cookie2state *c2s;
2197
2198             for (c2s = c2stbl; c2s->cookie != NULL; c2s++)
2199             {
2200                 if (strcmp(c2s->cookie, cp) == 0)
2201                 {
2202                     return c2s->state;
2203                 }
2204             }
2205             return INSTATE_SEARCH_PREFIX0;
2206             /* end of decodecookie() */
2207         }
2208
2209

```

```
2211 static void
2212 sigusr1_handler(int sigvaluel)
2213 {
2214     if (debugmode)
2215     {
2216         rbe_log_stats(0, "\nUSR1 (ATTN) signal received!\n");
2217     }
2218
2219     if (0 < xcpiogen_pid)
2220     {
2221         (void)kill(xcpiogen_pid, SIGTERM);
2222     }
2223
2224     ++attn_ec;
2225     /* end of sigusr1_handler() */
```

```
2227 static void
2228 sigterm_handler(int sigvaluel)
2229 {
2230     if (debugmode)
2231     {
2232         rbe_log_stats(0, "\nTERM signal received!\n");
2233     }
2234
2235     rbe_close_logs(logging_channel);
2236
2237     /* ESS workaround: Make a call to CDL_exit
2238     if symmpath workitem and CDL_exit has
2239     not already been called. This is to clean
2240     up sockets since we have made a call to
2241     CDL_noatexit, we have to make sure we
2242     clean up */
2243
2244     if ( (1==is_symmpath) && (0==sp_cdl_exitdone) )
2245     {
2246         CDL_exit();
2247         sp_cdl_exitdone = 1;
2248     }
2249
2250     signal(SIGTERM, SIG_DFL); /* restore default (terminate) action */
2251     (void)kill(getpid(), SIGTERM);
2252
2253     /* end of sigattn_handler() */
```

```

2256  /*
2257  * Invoked when running recover in debugging mode.
2258  * This execs a separate a.out file to implement the auxproc, so that
2259  * breakpoints can be set in the recover code without affecting
2260  * the auxproc code.
2261  */
2262
2263 static void
2264 z_exec_separate_auxproc(struct auxproc_context *exp,
2265                          char *pathname)
2266 {
2267     char procnum_str[32];
2268     char r_fd_str[32];
2269     char w_fd_str[32];
2270     char r_bulk_fd_str[32];
2271     char dbgmodestr[32];
2272     char *argv0;
2273
2274     /*
2275     * NOTE: This name is "special", the recover main()
2276     * function looks for it to know when it's being invoked
2277     * to perform auxproc processing.
2278     */
2279     argv0 = "ebr_auxproc";
2280
2281     (void) sprintf(procnum_str, "%d", exp->ap_my_auxnum);
2282     (void) sprintf(r_fd_str, "%d", exp->ap_r_fd);
2283     (void) sprintf(w_fd_str, "%d", exp->ap_w_fd);
2284     (void) sprintf(r_bulk_fd_str, "%d", exp->ap_r_bulk_fd);
2285     (void) sprintf(dbgmodestr, "%d", exp->ap_r_bulk_fd);
2286     (void) sprintf(dbgmodestr, "%d", debugmode);
2287     (void) execlp(pathname,
2288                  argv0,
2289                  /* prog to execute */
2290                  /* argv 0 */
2291                  procnum_str,
2292                  /* argv 1 */
2293                  r_fd_str,
2294                  /* argv 2 */
2295                  w_fd_str,
2296                  /* argv 3 */
2297                  r_bulk_fd_str,
2298                  /* argv 4 */
2299                  dbgmodestr,
2300                  /* argv 5 */
2301                  /* end-of-args */
2302                  (char *)0);
2303
2304     /*
2305     * if we get here, the exec did not go. Caller will bomb for us.
2306     */
2307 }
2308
2309 /* end of z_exec_separate_auxproc() */

```

```

2303  /*
2304  * Use these functions with looprw to build
2305  * a loop read (or loop write) that ignores EINTR.
2306  */
2307
2308 static int
2309 read_CD_L_no_eintr(int fd,
2310                   char *buf,
2311                   int nbytes)
2312 {
2313     int r;
2314
2315     do
2316     {
2317         errno = 0;
2318         r = CD_L_read(fd, buf, (uint_t)nbytes, 0);
2319         while (r == -1 && errno == EINTR);
2320     } while (r != 0);
2321
2322     return r;
2323 }
2324
2325 /* end of read_CD_L_no_eintr() */

```

```
2324 static int
2325 write_CDL_no_eintr(int fd,
2326                    char *buf,
2327                    int nbytes)
2328 {
2329     int r;
2330
2331     do
2332     {
2333         errno = 0;
2334         r = CDL_write(fd, buf, (uint_t)nbytes, 0);
2335         while (r == -1 && errno == EINTR);
2336     }
2337     return r;
2338 } /* end of write_CDL_no_eintr() */
```

```
2342 /*
2343  * Wait, interruptibly,
2344  * for at least one byte to become available on fd.
2345  * Returns 0 if at least one byte is available.
2346  * Returns -1 for any type of failure, including wait interruption.
2347  * Sets errno appropriate when -1 is returned.
2348  */
2349 int
2350 fd_avail_1_wait_intr(int fd)
2351 {
2352     esl_fset_ty rdbits;
2353
2354     E_FD_ZERO(&rdbits);
2355     E_FD_SET(fd, &rdbits);
2356
2357     /*
2358      * Don't need to examine rdbits after select, since only one
2359      * fd is in the set -- therefore return value can be computed
2360      * directly from select return value.
2361      */
2362     if (esl_select(E_FD_SETSIZE, &rdbits, NULL, NULL, NULL) == -1)
2363     {
2364         return -1;
2365     }
2366     return 0;
2367 } /* end of fd_avail_1_wait_intr() */
```



```

2372  /*
2373  * Test, interruptibly,
2374  *   for at least one byte to become available on fd.
2375  * Returns 1 if at least one byte is available.
2376  * Returns -1 for any type of failure, including test interruption.
2377  * Return 0 if no data is available.
2378  * Sets errno appropriate when -1 is returned.
2379  */
2380
2382  int
2383  fd_avail_test_intr(int fd)
2384  {
2386      int retStatus;
2387      struct pollfd fd_test;
2389      fd_test.fd = fd;
2390      fd_test.events = POLLIN;
2391      fd_test.revents = 0;
2394      /*
2395       * Don't need to examine rdbits after select, since only one
2396       * fd is in the set -- therefore return value can be computed
2397       * directly from select return value.
2398       */
2400      if ((retStatus = CDL_poll_read(&fd_test, 1, 0)) == -1)
2401      {
2402          /* ERROR encountered */
2403          return -1;
2404      }
2405      return retStatus;
2407  } /* end of fd_avail_test_intr() */

```

```

2409  /*
2410  * Test, for at least one byte to become available on fd.
2411  * Returns 1 if at least one byte is available.
2412  * Returns -1 for any type of failure other than EINTR.
2413  * Return 0 if no data is available.
2414  * Sets errno appropriate when -1 is returned.
2415  */
2417
2419  int fd_avail_test1(int fd)
2420  {
2421      int retStatus;
2422      struct pollfd fd_test;
2424      fd_test.fd = fd;
2425      fd_test.events = POLLIN;
2426      fd_test.revents = 0;
2429      /*
2430       * Don't need to examine rdbits after select, since only one
2431       * fd is in the set -- therefore return value can be computed
2432       * directly from select return value.
2433       */
2435      while ((-1 == (retStatus = CDL_poll_read(&fd_test, 1, 0))) &&
2436             (EINTR == errno))
2437      {
2438          ;
2439      }
2440      return retStatus;
2441  }

```

Thu Jan 03 12:25:21 2008		ebr_direct_rcmd		Page 105 of 144	
2443		#define EBR_FORK fork /* portability definition */	2503 1	if ((pid = EBR_FORK()) == 0)	
2445		static int ebr_direct_rcmd_fds[2] = { -1, -1 };	2504 2	{	
2447		/*	2505 2	/*	
2448		* Function to do a direct fork()/exec()	2506 2	* child processing goes here while parent is stopped (
2449		* client programs.	2507 2	* first, setup stdio to work using the pipes	
2450		Returns ptr to in & out fds for stdi/o with client process.	2508 2	*/	
2451		Returns ptr to fds if successful, NULL if an error. stderr from new	2510 2	if ((close(0)<0) (close(1)<0) (close(2)<0))	
2452		* process is returned on *fd2p.	2511 3	{	
2453		Should only be called when client == server.	2512 3	WriteFmtStringMsg(cxp->ap_w_prog_fd,	
2454		*/	2513 3	EDMREPROGMSG_AUXPROC_ERROR, 0,	
2455		int *	2514 3	unable to close std{	
2456		ubr_direct_rcmd(char	2515 3	in out err) for forked child\n");	
2457		ushort_t	2516 2	_exit(-1);	
2458		struct auxproc_context *cxp,	2517 2	if ((dup(x[0])<0) (dup(y[1])<0) (dup(z[1])<0))	
2459		char	2518 3	{	
2460		char	2519 3	WriteFmtStringMsg(
2461		int	2520 3	cxp->ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0,	
2462 1		{	2521 3	unable to dup pipe ends for forked	
2463 1		int pid;	2522 2	child\n");	
2464 1		int x[2];	2523 2	_exit(-1);	
2465 1		int y[2];	2524 2	if ((close(x[0])<0) (close(y[0])<0) (close(z[0])<0)	
2466 1		int z[2];	2525 3	{	
2467 1		int fd_w;	2527 3	WriteFmtStringMsg(
2468 1		int fd_r;	2528 3	cxp->ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0,	
2469 1		int fd_err;	2529 3	unable to close pipe ends for forked	
2470 1		char *client_home = NULL;	2530 2	child\n");	
2471 1		char *p;	2531 2	_exit(-1);	
2472 1		struct passwd *pw;	2532 2	/*	
2473 1		struct sigaction	2533 2	* At this moment the process has a real UID of the user	
2474 1		old_act;	2534 2	executing	
2475 1		int	2535 2	* [x]ebrecover and an effective UID of root (
2476 1			2536 2	since [x]ebrecover has	
2477 1			2537 2	* the setuid bit set and is owned by root).	
2478 1		/*	2538 2	* The first problem is the access(
2479 1		* now open pipes that will become daemons stdin, stdout, stderr	2539 2	* of the process. So if the user is not root or ebadmin,	
2480 1		* remember that x[0] is for read and x[1] is for write	2540 2	the access	
2481 1		*/	2541 2	* function will fail. This causes the recovery to fail.	
2482 1			2542 2	thus	
2483 1		if ((-1 == pipe(x)) (-1 == pipe(y)) (-1 == pipe(z)))	2543 2	* problem 1.	
2484 2		{	2544 2	* The second problem is System V Unix does not pass the	
2485 2		WriteFmtStringMsg(2545 2	effective UID	
2486 2		cxp->ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0,	2546 2	* to processes doing an exec of a shell.	
2487 2		Unable to create a pipe\n");	2547 2	* is changed back to the real UID of the process. However,	
2488 1		return NULL;	2548 2	BSD	
2489 1		fd_w = x[1]; /* pipe for child's stdin */	2549 2	* implementation does maintain the effective UID of the	
2490 1		fd_r = y[0]; /* pipe for child's stdout */	2550 2	process.	
2491 1		fd_err = z[0]; /* pipe for child's stderr */	2551 2	* On System V,	
2492 1		/*		neither the real or effective UID will be root during	
2493 1		* Prepare sigaction parameters		* the exec of the shell script if the user executing	
2494 1		* ignore SIGCHLD during this		[x]ebrecover is	
2495 1		sigemptyset(&new_act.sa_mask);		* a non-root user. This process will not be able to perform	
2496 1		new_act.sa_handler = SIG_IGN;		*/bin/sh -c.	
2497 1		new_act.sa_flags = E_SA_RESTRT;		/<nodename>/bin/startrec because this processes real	
2498 1		istat = sigaction(SIGCHLD, &new_act, &old_act);			
2499 1					
2500 1					
2501 1					
Thu Jan 03 12:25:21 2008		ebr_direct_rcmd		Page 106 of 144	

Page 107 of 144	ebr_direct_rcmd	Thu Jan 03 12:25:21 2008
2551 2	<i>* and effective UID (which are now the same) does not have</i>	
2552 2	<i>* permissions to the directory ~ebadmin.</i>	
2553 2	<i>* The recxpio process will be reading the username of the</i>	
2554 2	<i>* process from the standard input and setting the real UID to</i>	
2555 2	<i>* UID of that user so it will do no harm to set the real UID</i>	
2556 2	<i>* at this point.</i>	
2557 2	<i>* It will be changed back to the actual user as one</i>	
2558 2	<i>* of the first things in recxpio.</i>	
2559 2	<i>*/</i>	
2560 2	<i>(void)setuid(geteuid());</i>	
2561 2	<i>*/</i>	
2562 2	<i>* now chdir to the client code home location</i>	
2563 2	<i>*/</i>	
2564 2	<i>if (NULL == (pw = getpwnam(</i>	
2565 2	<i>remuser))) /* lookup client home dir */</i>	
2566 2	<i>{</i>	
2567 3	<i> p = "Can't get home directory";</i>	
2568 3	<i> goto default_home;</i>	
2569 3	<i> }</i>	
2570 3	<i> else</i>	
2571 3	<i> {</i>	
2572 2	<i> if (</i>	
2573 2	<i> client_home = pw->pw_dir) != NULL) /* if there is a pointer */</i>	
2574 3	<i> {</i>	
2575 3	<i> if (0 != chdir(client_home)) /* cd to directory */</i>	
2576 4	<i> {</i>	
2577 5	<i> p = "Can't chdir to home";</i>	
2578 5	<i> goto default_home;</i>	
2579 5	<i> }</i>	
2580 5	<i> else</i>	
2581 4	<i> {</i>	
2582 3	<i> p = "Can't figure out home dir";</i>	
2583 3	<i> }</i>	
2584 4	<i> }</i>	
2585 4	<i> }</i>	
2586 4	<i> p = "Can't figure out home dir";</i>	
2587 4	<i> writeFmtStringMsg(cxp->ap_w_prog_fd,</i>	
2588 4	<i> EDMREPROGMSG_AUXPROC_WARNING, 0,</i>	
2589 4	<i> "%s";</i>	
2590 4	<i> /* %s\ for user \"%s\", defaulting to /usr/epoch/EB\n",</i>	
2591 4	<i> p, (</i>	
2592 3	<i> client_home == NULL) ? "??": client_home, remuser);</i>	
2593 2	<i> client_home = "/usr/epoch/EB/CLIENT_HOME";</i>	
2594 2	<i> }</i>	
2595 2	<i> /*</i>	
2596 2	<i> client_home now points to the home dir of the client</i>	
2597 2	<i> software</i>	
2598 2	<i> /*</i>	
2599 2	<i> Make sure that the home directory that we finally ended</i>	
2600 2	<i> up with</i>	
2601 2	<i> /*</i>	
2602 2	<i> is really there.</i>	
2603 2	<i> /*</i>	
2604 2	<i> Yea. I know.</i>	
	<i> I may have already done this if I found a passwd</i>	
	<i> entry for the client backup username and it had a home</i>	
	<i> directory.</i>	

Page 108 of 144	ebr_direct_rcmd	Thu Jan 03 12:25:21 2008
2605 2	<i>* However,</i>	
2606 2	<i> I have to do it again just in case I came from another</i>	
2607 2	<i> patch.</i>	
2608 2	<i> I did not really feel like trying to fix up the spaghetti</i>	
	<i> code. I'm feeling a little lazy today...</i>	
2610 2	<i> if (0 != chdir(client_home))</i>	
2611 3	<i> {</i>	
2612 3	<i> writeFmtStringMsg(</i>	
2613 3	<i> cxp->ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0,</i>	
2614 3	<i> "Server backup directory \"%s\" not found</i>	
2615 3	<i> or not searchable\n",</i>	
2616 2	<i> client_home);</i>	
2617 2	<i> }_exit(-1);</i>	
2618 2	<i> (void)exec1("/bin/sh", "sh", "-c", cmd, 0);</i>	
2619 2	<i> writeFmtStringMsg(</i>	
2620 2	<i> cxp->ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0,</i>	
2621 2	<i> "Unable to exec1 /bin/sh sh -c %s\n", cmd);</i>	
2622 1	<i> }_exit(127);</i>	
2623 1	<i> }</i>	
2624 1	<i> /*</i>	
2625 1	<i> parent code resumes here</i>	
2626 1	<i> /*</i>	
2627 1	<i> Close the parent's copies of the child-ends of the pipes</i>	
2628 1	<i> /*</i>	
2629 1	<i> close(x[0]);</i>	<i>/* pipe for child's stdin */</i>
2630 1	<i> close(y[1]);</i>	<i>/* pipe for child's stdout */</i>
2631 1	<i> close(z[1]);</i>	<i>/* pipe for child's stderr */</i>
2632 1	<i> /*</i>	
2633 1	<i> check for fork() failure</i>	
2634 1	<i> /*</i>	
2635 1	<i> if (-1 == pid)</i>	
2636 1	<i> {</i>	
2637 1	<i> writeFmtStringMsg(</i>	
2638 1	<i> cxp->ap_w_prog_fd, EDMREPROGMSG_AUXPROC_ERROR, 0,</i>	
2639 1	<i> "fork() failed\n");</i>	
2640 1	<i> }</i>	
2641 2	<i> close(x[1]);</i>	<i>/* pipe for child's stdin */</i>
2642 2	<i> close(y[0]);</i>	<i>/* pipe for child's stdout */</i>
2643 2	<i> close(z[0]);</i>	<i>/* pipe for child's stderr */</i>
2644 2	<i> if (istat == 0)</i>	
2645 2	<i> {</i>	
2646 2	<i> sigaction(SIGCHLD, {old_act, NULL};</i>	
2647 2	<i> return NULL;</i>	
2648 2	<i> }</i>	
2649 3	<i> }</i>	
2650 3	<i> /*</i>	
2651 2	<i> now fill in fd's to be returned</i>	
2652 2	<i> /*</i>	
2653 1	<i> _ebr_direct_rcmd_fds[0] = fd_w;</i>	<i>/* pipe for child's stdin */</i>
2654 1	<i> _ebr_direct_rcmd_fds[1] = fd_r;</i>	<i>/* pipe for child's stdout */</i>
2655 1	<i> if (NULL != fd2p)</i>	
2656 1	<i> {</i>	
2657 1	<i> *fd2p = fd_err;</i>	
2658 1	<i> }</i>	

```

2666 1      if (istat == 0)
2667 2      {
2668 2          sigaction(SIGCHLD, &old_act, NULL);
2669 1      }
2671 1      return ebr_direct_rcmd_fds;
2672 1      } /* end of ebr_direct_rcmd() */

```

```

2674      /*
2675      * routine to search the clients_installed file for a clients
2676      * method. Returns a pointer to the connection method string if
2677      * NULL if the client entry could not be found.
2678      */
2680      static char *
2681      rb_getmethod(register char *host,
2682                  uint_t *concurrency,
2683                  uint_t *client_type)
2684      {
2685          FILE
2686          *file_ptr;
2687          FILE
2688          *lock_ptr;
2689          FFDB_HOLDER
2690          *incient_holder;
2691          ffdb_incident_ty
2692          *installed_client;
2693          errno;
2694          static char
2695          rb_ci_lastmethod[32] = "";

```

```

2696      #ifdef PARAM_CHECK
2697      if (NULL == host)
2698      {
2699          rbe_internal_error(RBRECOVER_MKERR(
2700              EINVAL), null_arg, "host name");
2701          return NULL;
2702      }
2703      #endif /* PARAM_CHECK */

```

```

2704      /*
2705      * get pathname to file
2706      */
2707      if (ffdb_incident_path == NULL) /* if file name needs to be
2708                                      prepared */
2709      {
2710          if (ffdb_incident_init(FFDB_INIT_DEFAULT_PATH) != 0)
2711          {
2712              return NULL;
2713          }

```

```

2714      /*
2715      * request was not in cache -- go load it
2716      */

```

```

2717      if ((lock_ptr = ffdb_lock_file(ffdb_incident_path,
2718                                     "restore scanning
2719                                     clients_installed file",
2720                                     1, NULL, 0)) == NULL)

```

```

2721      {
2722          (void)rbe_user_error(
2723              ffdb_errnum, "Unable to obtain lock on \"%s\" for scanning",
2724              ffdb_incident_path);
2725          return NULL;
2726      }

```

```

2727      if ((file_ptr = ffdb_open_file(ffdb_incident_path)) == NULL)
2728      {
2729          (void)rbe_user_error(
2730              ffdb_errnum, "Unable to open file \"%s\" for scanning",
2731              ffdb_incident_path);
2732          ffdb_unlock_file(lock_ptr);
2733          return NULL;
2734      }

```

```

2733 1      while ((inclient_holder = ffdb_read_record(
2734 2          file_ptr, &errnum)) != NULL)
2735 2      {
2736 2          installed_client = (ffdb_inclient_ty *)inclient_holder->drec;
2737 3          if (errnum == 0)
2738 3              if (ebcd_same_host(
2739 4                  host, installed_client->hostname)) /* a match? */
2740 4                  {
2741 5                      if (concurrency != NULL)
2742 5                          if (0 == installed_client->concurrency)
2743 6                              {
2744 6                                  *concurrency = 32767;
2745 5                              }
2746 5                              else
2747 6                                  {
2748 6                                      *concurrency = installed_client->concurrency;
2749 5                                  }
2750 4                              }
2751 4                              if (client_type != NULL)
2752 5                                  {
2753 5                                      *client_type = installed_client->platform;
2754 4                                  }
2755 4                                  strcpy(rb_ci_lastmethod, installed_client->method);
2756 4                                  (void)ffdb_close_file(file_ptr);
2757 4                                  ffdb_unlock_file(lock_ptr);
2758 4                                  ffdb_destroy_holder(inclient_holder);
2759 4                                  return rb_ci_lastmethod; /* return allocated string */
2760 3                                  }
2761 2                                  ffdb_destroy_holder(inclient_holder);
2762 2                                  }
2763 1                                  }
2765 1                                  /* if we get here -- did not find host in clients_installed file
2766 1                                  */
2767 1                                  */
2769 1                                  (void)ffdb_close_file(file_ptr);
2770 1                                  ffdb_unlock_file(lock_ptr);
2771 1                                  return NULL;
2772 1                                  /* end of rb_getmethod() */

```

```

2775 1      /*
2776 1      * Reads remote file descriptor for reads of stderr
2777 1      * of the remote cmd. Read is done until nothing is left
2778 1      * then returns with or without exit code.
2779 1      * The error/warning messages come first and a dealt with,
2780 1      * if the exit code of the remote cmd is returned.
2781 1      */
2782 1      (Inp) int fd -- The file descriptor to wait on and read.
2783 1      * (Out) int *exitp -- The Exit code.
2784 1      * (Inp) char *remhostname -- The remote client name.
2785 1      * (Inp) boolean_ty first_call -- Is this the first call ??
2786 1      * (I/O) enum input_states *state_ptr
2787 1      * (I/O) enum input_states *next_state_ptr
2788 1      * (I/O) boolean_ty *skipping_leading_whitespace.
2789 1      * (I/O) int *parsepos the parse position of the messages.
2790 1      * (I/O) int *msgpos the position of the logging message.
2791 1      *
2792 1      * None of the I/O vars need to be initialized for the first call.
2793 1      *
2794 1      * Returns: boolean_ty
2795 1      * False: The remote command has not exited.
2796 1      * True: The remote command has exited.
2797 1      */
2798 1      static boolean_ty
2799 1      parse_remote_stder_info2(int prog_fd,
2800 1          int remote_fd,
2801 1          int *exitp,
2802 1          char *remhostname,
2803 1          boolean_ty first_call,
2804 1          enum input_states *state_ptr,
2805 1          enum input_states *next_state_ptr,
2806 1          boolean_ty *skipping_leading_whitespace,
2807 1          int *parsepos,
2808 1          int *msgpos)
2809 1      {
2810 1          enum input_states previous_state;
2811 1          static char cookiebuf[RMPD_MAX_COOKIE_LEN+1];
2812 1          int profailed = 0;
2813 1          int done = 0;
2814 1          char c;
2815 1          int n;
2816 1          #define MSGBUFLLEN 260
2817 1          static char msglogBuf[MSGBUFLLEN + 1];
2818 1          static boolean_ty ret_status = FALSE;
2819 1          static int temp_exit_status;
2820 1          int write_prog = 1;
2821 1          int write_prog = 1;
2822 1          *exitp = 0; /* Initialize to success. */
2823 1          if (TRUE == first_call)
2824 1          {
2825 1              /* First time processing. */
2826 1              if (debugmode)
2827 1                  rbe_log_stats(
2828 2                      0, "AUXPROC: Processing stderr info from fd %d\n",
2829 2                      remote_fd);
2830 1              }
2831 1              *state_ptr = INSTSTATE_NG;
2832 2              *next_state_ptr = INSTSTATE_SEARCH_PREFIX0;
2833 2              *parsepos = 0;
2834 2              *msgpos = 0;
2835 2              *parsepos = 0;
2836 2              *msgpos = 0;
2837 2

```

```

2838 2      *exitp = 0;
2839 2      *msgPos = 0;
2840 2      temp_exit_status = 0;
2841 2      *skipping_leading_whitespace = FALSE;
2842 2      n = 0;
2843 1      }
2844 1      else
2845 2      {
2846 2          n = *parsePos;
2847 1      }
2849 1      while (!protofailed && !done)
2850 2      {
2851 2          int start_ec;
2852 2          int r;

2855 2      /*
2856 2      * Wait for the next character from the remote
2857 2      * stream. Ignore interrupts, unless
2858 2      * they were due to the attention signal.
2859 2      */

2861 2      start_ec = atn_ec;
2862 2      do
2863 3      {
2864 3          r = fd_await_test_intr(remote_fd);
2865 3          if (0 == r)
2866 4          {
2867 4              return ret_status;
2868 3          }
2869 2          while (r == -1 && errno == EINTR && atn_ec == start_ec);

2871 2      /*
2872 2      * now that there is a character, read it
2873 2      */
2874 2      c = 0;

2876 2      if (r != -1)
2877 3      {
2878 3          do
2879 4          {
2880 4              r = CDL_read(remote_fd, &c, 1, 0);
2881 3              while((-1 == r) && (errno == EINTR) && (
                atn_ec == start_ec));
2882 2          }
2884 2          if (r != 1)
2885 3          {
2886 3              *exitp = SPEXIT_REMOTE_STDBRR_FAIL;
2887 3              protofailed = 1;
2888 3              ret_status = TRUE;
2889 3              break;
2890 2          }

2892 2          previous_state = *state_ptr;
2893 2          *state_ptr = *next_state_ptr;

2895 2          switch (*state_ptr)
2896 3          {
2897 3              case INSTSTATE_SEARCH_PREFIX0:
2898 3                  if (debugmode)
2899 4                  {
2900 4                      rbe_log_stats(0, "AUXPROC: (
                %c) INSTSTATE_SEARCH_PREFIX0\n", c);
2901 3                  }

```

```

2903 3          if (c == REMFD_MAGIC_PREFIX(0))
2904 4          {
2905 4              *next_state_ptr = INSTSTATE_SEARCH_PREFIXN;
2906 4              n = 1;
2907 4              *parsePos = 1;
2908 3          }
2909 3          break;

2911 3      case INSTSTATE_SEARCH_PREFIXN:
2912 3          if (debugmode)
2913 4          {
2914 4              rbe_log_stats(0, "AUXPROC: (
                %c) INSTSTATE_SEARCH_PREFIXN\n", c);
2915 3          }

2917 3          if (c != REMFD_MAGIC_PREFIX(n))
2918 4          {
2919 4              *next_state_ptr = INSTSTATE_SEARCH_PREFIX0;
2920 3          }
2921 3          else
2922 4          {
2923 4              n++;
2924 4              (*parsePos)++;
2925 4              if (n == REMFD_MAGIC_LENGTH)
2926 5              {
2927 5                  *next_state_ptr = INSTSTATE_GATHER_COOKIE;
2928 5                  n = 0;
2929 5                  *parsePos = 0;
2930 4              }
2931 3          }
2932 3          break;

2934 3      case INSTSTATE_GATHER_COOKIE:
2935 3          if (debugmode)
2936 4          {
2937 4              rbe_log_stats(0, "AUXPROC: (
                %c) INSTSTATE_GATHER_COOKIE\n", c);
2938 3          }

2940 3          if (c == REMFD_MAGIC_SUFFIX(0))
2941 4          {
2942 4              /*
2943 4              * Got the cookie
2944 4              */

2946 4              cookiebuf[n] = '\0';
2947 4              *next_state_ptr = INSTSTATE_SEARCH_SUFFIXN;
2948 4              n = 1;
2949 4              *parsePos = 1;
2950 3          }
2951 3          else if (strchr(REMFD_COOKIE_CHARS, c) == NULL)
2952 4          {
2953 4              /*
2954 4              * We found a valid prefix (else we would
2955 4              * not be here) but the cookie contains an
2956 4              * illegal character. Should not happen;
2957 4              * there has been some protocol failure.
2958 4              */

2960 4              *exitp = SPEXIT_REMOTE_STDBRR_FAIL;
2961 4              ret_status = TRUE;
2962 4              protofailed = 1;
2963 3          }
2964 3          else if (n == REMFD_MAX_COOKIE_LEN)
2965 4          {

```

```
2966 4 /*
2967 4  * Should have seen the SUFFIX[0] by now.
2968 4 */
2970 4 *exitp = SPEXIT_REMOTE_STDPERR_PROTOCOL;
2971 4 ret_status = TRUE;
2972 4 protfailed = 1;
2973 3 }
2974 3 else
2975 4 {
2976 4     /*
2977 4     * another cookie character.
2978 4     */
2980 4     cookiebuf[n] = c;
2981 4     n++;
2982 4     (*parsePos)++;
2983 3 }
2984 3 break;

2986 3 case INSTSTATE_SEARCH_SUFFIXN:
2987 3     if (debugmode)
2988 4         rbe_log_stats(0, "AUXPROC: (
2989 4         %c) INSTSTATE_SEARCH_SUFFIXN\n", c);
2990 3 }
2992 3 if (c != REMFD_MAGIC_SUFFIX[n])
2993 4 {
2994 4     *exitp = SPEXIT_REMOTE_STDPERR_PROTOCOL;
2995 4     ret_status = TRUE;
2996 4     protfailed = 1;
2997 3 }
2998 3 else
2999 4 {
3000 4     n++;
3001 4     (*parsePos)++;
3002 4     if (n == REMFD_MAGIC_LENGTH)
3003 5     {
3004 5         *next_state_ptr = INSTSTATE_NEWLINE;
3005 4     }
3006 3 }
3007 3 break;

3009 3 case INSTSTATE_GATHER_STATUS:
3010 3     if (debugmode)
3011 4     {
3012 4         rbe_log_stats(0, "AUXPROC: (
3013 4         %c) INSTSTATE_GATHER_STATUS\n", c);
3014 3     }
3015 3     if (isdigit(c))
3016 4     {
3017 4         char xx[2];
3018 4         xx[0] = c;
3019 4         xx[1] = '\0';
3020 4         temp_exit_status *= 10;
3021 4         temp_exit_status += atoi(xx);
3022 4         n++;
3023 4         (*parsePos)++;
3024 4         *skipping_leading_whitespace = FALSE;
3025 4     }
3026 4     else if (c == '\n')
3027 3 }
3028 3 }
3029 3
```

```
3030 4 {
3031 4     *exitp = temp_exit_status;
3032 4     ret_status = TRUE;
3033 4     done = 1;
3034 3 }
3035 3 else if (c == ' ' &&
3036 3     (n == 0) || *skipping_leading_whitespace )
3037 4 {
3038 4     *skipping_leading_whitespace = TRUE;
3039 4 }
3040 4 /*
3041 4     * no-op, skip leading whitespace
3042 4     */
3043 3 }
3044 3 else
3045 4 {
3046 4     *exitp = SPEXIT_REMOTE_STDPERR_PROTOCOL;
3047 4     ret_status = TRUE;
3048 4     protfailed = 1;
3049 3 }
3050 3 break;

3052 3 case INSTSTATE_COPY_TO_STDOUT:
3053 3 {
3054 3     /*
3055 3     * Eventually, the protocol should include
3056 3     * an explicit length for this state, so
3057 3     * we can do large read/writes and so we will not
3058 3     * be vulnerable to confusion based on gunk being
3059 3     * copied to stdout.
3060 3     * For now, just shove the characters at stdout
3061 3     * and stop as soon as we fix PREFIX[0]
3062 3     */
3063 3     if (c != REMFD_MAGIC_PREFIX[0])
3064 4     {
3065 4         /*
3066 4         * build a buffer to write to the restore log
3067 4         */
3068 4         if (('\'\' == c) || ((*msgPos) >= MSGBUFLEN))
3069 5         {
3070 5             msgLogBuf[*msgPos] = '\0';
3071 5             (void) rbe_log_stats(0,
3072 5                 "s: %s",
3073 5                 remhostname,
3074 5                 msgLogBuf);
3075 5             if (write_prog)
3076 6             {
3077 6                 writeStringMsg(prog_fd,
3078 6                     EDMREPROGMSG_RESTORE_RCMD_UNKNOWN,
3079 6                     0,
3080 6                     msgLogBuf);
3081 6             }
3082 6             *msgPos = 0;
3083 5         }
3084 5         memset(msgLogBuf, 0, MSGBUFLEN + 1);
3085 5         *msgPos = 0;
3086 5     }
3087 5     else
3088 4     {
3089 4         msgLogBuf[*msgPos] = c;
3090 4         (*msgPos)++;
3091 5     }
3092 5 }
3093 4
```

```
3095 3      }
3096 3      else
3097 4      {
3098 4          n = 1;
3099 4          *parsePos = 1;
3100 4          next_state_ptr = INSTSTATE_SEARCH_PREFIXIN;
3101 3      }
3102 3      break;
3104 3      case INSTSTATE_NEWLINE:
3105 3          if (c == '\n')
3106 4          {
3107 5              if (previous_state == INSTSTATE_SEARCH_SUFFIXIN) {
3108 5                  *next_state_ptr = decodecookie(cookiebuf);
3110 5                  memset(cookiebuf, 0, REMFD_MAX_COOKIE_LEN+1);
3112 5                  n = 0;
3113 5                  *parsePos = 0;
3114 4              }
3115 4              else
3116 5              {
3117 5                  *next_state_ptr = INSTSTATE_SEARCH_PREFIXIN;
3118 4              }
3119 3          }
3120 3          break;
3121 2          /* end of switch */
3122 1          /* end of while loop */
3124 1      }
3126 1      return ret_status;
          /* end of parse_remote_slderr_info2() */
```

```
3131 1      /*
3132 1      * ForwardXcpiogenProgress()
3133 1      * REP: Restore Engine Process, AP: auxproc, XC: xcpiogen.
3134 1      * This routine read progress messages from XC and forwards
3135 1      * then to REP. This routine should not block on a read().
3136 1      */
3137 1      * Args:
3138 1      * int xcpiogen_prog_fd -- the progress channel from XC to AP.
3139 1      * int restore_engine_prog_fd -- the progress channel from AP to REP.
3140 1      * Returns:
3141 1      * int -- the number of forwarded messages. -1 for an error.
3143 1      * Side Effects:
3144 1      * Read and writes file descriptors.
3145 1      */
3146 1      static int
3147 1      ForwardXcpiogenProgress(int xcpiogen_prog_fd,
3148 1      int restore_engine_prog_fd,
3149 1      boolean_ty *zero_byte_read)
3150 1      {
3151 1          int bytes_skipped = 0;
3152 1          int read_ret, write_ret;
3153 1          int fd_test;
3154 1          int msg_count = 0;
3155 1          prog_chan_msg xcpiogen_msg;
3156 1          int debug_prog = 0;
3157 1          int write_prog = 1;
3158 1          int write_byte_read = 1;
3160 1          *zero_byte_read = FALSE;
3162 1          memset(&xcpiogen_msg, 0, sizeof(prog_chan_msg));
3164 1          while(1 == (fd_test = fd_avail_test1(xcpiogen_prog_fd)))
3165 2          {
3166 2              read_ret = ReadMsg(
3168 2                  xcpiogen_prog_fd, &xcpiogen_msg, &bytes_skipped);
3169 2              if(0 != bytes_skipped)
3170 2              {
3171 2                  (void) rbe_log_stats(0,
3172 2                      "The progress channel from xcpiogen had %d
3173 2                      extraneous bytes.");
3174 2              }
3175 2              if(0 == read_ret)
3176 2              {
3177 2                  *zero_byte_read = TRUE;
3178 2                  break;
3179 2              }
3180 2              else if(-1 == read_ret)
3181 2              {
3182 2                  break;
3183 2              }
3184 2              if(debug_prog)
3185 2              {
3186 2                  (void) rbe_log_stats(0,
3187 2                      "%s: %s",
3188 2                      "debug progress",
3189 2                      DEBUG_PROG_MESSAGE_PEAK (&xcpiogen_msg));
3190 2              }
3191 2          }
```



```

3191 3         }
3192 2         write_ret = WriteMsg(restore_engine_prog_fd, &xcpiogen_msg);
3193 2     }
3194 2     msg_count++;
3196 2     free(xcpiogen_msg.pcm_body);
3197 1     memset(&xcpiogen_msg, 0, sizeof(prog_chan_msg));
3198 1     if(-1 == fd_test)
3199 2     {
3200 2         (void) rbe_log_stats(RBRCOVER_MKERR(errno),
3201 2             "Encountered error testing xcpiogen file
3202 2             descriptor.");
3203 1     }
3204 1     return -1;
3205 1     }
3206 1     return msg_count;
3207 1     }

```

```

3207 1     /*
3208 1     * static int DemuxAuxChildren()
3209 1     * This function handles IPC communications between the following
3210 1     * processes: REP: Restore Engine Process, AP: auxproc, XC: xcpiogen,
3211 1     * and RC: remote command. The RC sends error and warning messages
3212 1     * to AP. When the RP is finished, RC's exit status is sent last.
3213 1     * The remote exit status indicates that the restore is finished.
3214 1     * RP error and warning messages are logged and forward on as
3215 1     * progress messages to RP. At the same time XC will send progress
3216 1     * information to AP, and AP will forward them to REP. XC does its
3217 1     * own logging.
3218 1     *
3219 1     * Args:
3220 1     * int progress_fd -- (In) this file descriptor is from AP to REP.
3221 1     * int remote_fd -- (In) this file descriptor is from RC to AP.
3222 1     * char *remote_programme -- (In) this the RP executable name.
3223 1     * int xcpiogen_fd -- (In) this file descriptor is from XC to AP.
3224 1     * int xcpiogen_pid -- (In) the pid for XC.
3225 1     * int *remote_exit -- (Out) the remote exit interpreted.
3226 1     *
3227 1     * REP: Restore Engine Process, AP: auxproc, XC: xcpiogen,
3228 1     * RC: remote command.
3229 1     *
3230 1     * Returns: ??
3231 1     *
3232 1     */
3233 1
3234 1     static int
3235 1     DemuxAuxChildren(int progress_fd,
3236 1                     int remote_fd,
3237 1                     char *remote_programme,
3238 1                     int xcpiogen_fd,
3239 1                     int xcpiogen_pid,
3240 1                     int *remote_exit)
3241 1     {
3242 1
3243 1         int remote_exitinfo;
3244 1         int forward_ret;
3245 1         boolean_ty remote_exitted = FALSE;
3246 1         struct pollfd monitor_fds[2];
3247 1         int poll_ret;
3248 1         boolean_ty remote_first_call = TRUE;
3249 1         int logging_channel = 0;
3250 1         int start_ec = atn_ec; /* atn_ec global */
3251 1         boolean_ty xcpiogen_zero_byte_reads;
3252 1         int number_fds;
3253 1         /* The below args are needed by and maintained by
3254 1         * parse_remote_stterr_info2(). No need to initialize
3255 1         * or test them, they simply maintain state information
3256 1         * for multiple invocations of parse_remote_stterr_info2().
3257 1         */
3258 1
3259 1         enum input_states remote_state_ptr;
3260 1         enum input_states remote_next_state_ptr;
3261 1         boolean_ty skip_whitespace;
3262 1         int parsePos;
3263 1         int MsgPos;
3264 1
3265 1         number_fds = 2;
3266 1         monitor_fds[0].fd = remote_fd;
3267 1         monitor_fds[0].events = POLLIN;
3268 1         monitor_fds[0].revents = 0;

```

```

3270 1      monitor_fds[1].fd = xcpiogen_fd;
3271 1      monitor_fds[1].events = POLLIN;
3272 1      monitor_fds[1].revents = 0;

3274 1      while((FALSE == remote_exitted) &&
3275 1          ((0 == (poll_ret = CDL_poll_read(
3276 1              (poll_ret > 0) ||
3277 1                  ((EINTR == errno) && (-1 == poll_ret) && (
3278 2                      start_ec != attn_ec))))
3279 2              {
3280 2                  if (0 == poll_ret)
3281 3                      /* timed out */
3282 3                      /* check for the attention signal and other periodic
3283 3                          checks. */
3284 4                      if (start_ec != attn_ec)
3285 4                          {
3286 4                              /* Timeout on read, AND we were canceled (SIGUSR1) --
3287 4                                  If this is an SSL socket read timeout, the client may
3288 4                                      have gone away without us knowing it,
3289 4                                          so break out of the loop
3290 4                                              without the remote exit status.
3291 4                                                  Do this for all remote
3292 4                                                      connection might be SSL,
3293 4                                                          since the xcpiogen to rexcpio
3294 4                                                              connection might be SSL,
3295 4                                                                  but not the remote_fd. In this case
3296 4                                                                      the remote process will not know that xcpiogen is
3297 4                                                                          gone.
3298 4                                                                              */
3299 4                                                                              *remote_exit = SPEXIT_REMOTE_NO_STATUS;
3300 4                                                                              /* leaving w/o status */
3301 3                                                                              }
3302 2                                                                              }
3303 2                                                                              else if (-1 == poll_ret)
3304 3                                                                              { /* eintr */
3305 3                                                                              /* log this here */
3306 3                                                                              /*rbe_log_stats(errno, "auxproc -- ");*/
3307 2                                                                              }
3308 2                                                                              else
3309 3                                                                              {
3310 3                                                                              if((POLLRDNORM & monitor_fds[1].revents) ||
3311 3                                                                              (POLLRDBAND & monitor_fds[1].revents)
3312 3                                                                              (POLLIN & monitor_fds[1].revents) ||
3313 3                                                                              (POLLRHUP & monitor_fds[1].revents))
3314 4                                                                              {
3315 4                                                                              /* Xcpiogen process is expected to send progress
3316 4                                                                              formatted
3317 4                                                                              messages,
3318 4                                                                              including warnings and errors along with restore
3319 4                                                                              progress over its stderr channel.
3320 4                                                                              We will check xcpiogen
3321 4                                                                              prior to the remote channel.
3322 4                                                                              This presumes that the remote
3323 4                                                                              channel determines whether we are done with the
3324 4                                                                              restore.

```

```

3320 4      /*
3321 4      monitor_fds[1].revents = 0;
3322 4      xcpiogen_zero_byte_reads = FALSE;
3323 4      forward_ret = ForwardXcpiogenProgress(
3324 4          xcpiogen_prog_fd,
3325 4          progress_fd,
3326 4          &xcpiogen_zero_byte_reads);
3327 4      if(TRUE == xcpiogen_zero_byte_reads)
3328 4      {
3329 5          /* Lets no longer wait on this fd,
3330 5          * if we get a zero byte read. The
3331 5          * second fd is for xcpiogen.
3332 5          */
3333 5          number_fds = 1;
3334 5          }
3335 4      }
3336 3      if((POLLRDNORM & monitor_fds[0].revents) ||
3337 3          (POLLRDBAND & monitor_fds[0].revents) ||
3338 3          (POLLIN & monitor_fds[0].revents) ||
3339 3          (POLLRHUP & monitor_fds[0].revents))
3340 3      {
3341 4          /* POLLRDBAND why do we need these conditions. */
3342 4          {
3343 4              /* Remote process is expected to send information back
3344 4                  stream. Read that information, parse it,
3345 4                  and if available
3346 4                  * send remote exit status back to parent.
3347 4                  If the exit status
3348 4                  * is read this loop is terminated.
3349 4                  */
3350 4                  monitor_fds[0].revents = 0;
3351 4                  remote_exitted = parse_remote_stderr_info2(progress_fd,
3352 4                                                                  remote_fd,
3353 4                                                                  &remote_exitinfo,
3354 4                                                                  remote_progname,
3355 4                                                                  remote_first_call,
3356 4                                                                  &remote_state_ptr,
3357 4                                                                  &remote_next_state_ptr,
3358 4                                                                  &skip_whitespace,
3359 4                                                                  &parsePos,
3360 4                                                                  &msgPos);
3361 4                  remote_first_call = FALSE;
3362 3                  }
3363 2                  }
3364 1                  }
3365 1                  /* Should we test xcpiogen for a last time before we return?
3366 1                  */
3367 1                  if(1 == fd_avail_test(xcpiogen_prog_fd))
3368 1                  {
3369 2                      xcpiogen_zero_byte_reads = FALSE;
3370 2                      forward_ret = ForwardXcpiogenProgress(xcpiogen_prog_fd,
3371 2                                                                  forward_ret);

```

progress_fd,

&xcpio_gen_zero_byte_reads);

}

if (-1 == poll_ret)

{ /* Not eintr */

rbe_log_stats(RBRECOVER_MKERR(errno),

"Auxproc failed to poll Children pids!");

return -1;

if(TRUE == remote_exitted)

{ *remote_exit = remote_exitinfo;

return 0;

}

} /* DemuxAuxChildren() */

```
1  /*
2      RSLauxsupp.c
3
4  * Copyright (C) 1998 by EMC Corp, Inc. All rights reserved.
5
6  *
7  * Table of Contents:
8  *-----
9  ** int looprw(int fd, char *buf, int nbytes, int (*func)());
10
11 */
12
13 #define _POSIX_SOURCE 1
14 #include <eb/eb_port.h>
15
16
17 #include <netdb.h>
18 #include <pwd.h>
19 #include <ctype.h>
20 #include <sys/wait.h>
21 #include <sys/types.h>
22 #include <util/esl_select.h>
23 #include <errno/esl_strerror.h>
24 #include <util/esl_limit.h>
25 #include <stdlib.h>
26
27
28 #include <ebconfig/rbconfig.h>
29 #include <epfddb/ffdb.h>
30 #include <epfddb/rbifclient.h>
31 #include <ebutil/ebsock_if.h>
32 #include <eb/rb_log.h>
33
34 #include <RSLinterns.h>
35 #include <RSLspexit.h>
36 #include <RSLremfd.h>
37 #include "RSLauxSupp.h"
38
39
40 static char *oops_msg = NULL;
41 static char oops_string[1024];
42
43
44 /*
45  * read() or write() a file descriptor, looping as
46  * necessary until all bytes are obtained. Useful for reading
47  * from pipes or sockets, or anything else that is permitted to
48  * return fewer bytes than you ask for.
49  */
50
51 int
52 looprw(int fd,
53        char *buf,
54        int nbytes,
55        int (*func)(int fd, char *buf, int nbytes))
56 {
57     int todo;
58
59     if (nbytes <= 0)
60         /* (in)sanity */
61         return nbytes;
62 }
63
64 /*
65  */
```

```
67     Each time try to 'func' the entire amount remaining.
68     * Exit the loop successfully when we've done it all.
69     * If 'func' returns an error, propagate the error.
70     * If 'func' returns EOF, return the amount done
71     * prior to the EOF.
72
73     todo = nbytes;
74     for (;;)
75     {
76         int r;
77
78         r = (*func)(fd, buf, todo);
79         if (r == todo)
80             /* all done */
81             return nbytes;
82         else if (r == 0)
83             /* EOF */
84             return nbytes - todo;
85         else if (r == -1)
86             /* error */
87             return -1;
88     }
89
90     buf += r;
91     todo -= r;
92
93     /*NOTREACHED*/
94     /* end of looprw() */
95 }
```

```

104  /*
105  * Use these functions with looprw to build
106  * a loop read (or loop write) that ignores EINTR.
107  */
108
109  int
110  read_no_eintr(int fd,
111               char *buf,
112               int nbytes)
113  {
114      int r;
115
116      do
117      {
118          errno = 0;
119          r = read(fd, buf, (uint_t)nbytes);
120      } while (r == -1 && errno == EINTR);
121
122      return r;
123  } /* end of read_no_eintr() */

```

```

125  int
126  write_no_eintr(int fd,
127                char *buf,
128                int nbytes)
129  {
130      int r;
131
132      do
133      {
134          errno = 0;
135          r = write(fd, buf, (uint_t)nbytes);
136      } while (r == -1 && errno == EINTR);
137
138      return r;
139  } /* end of write_no_eintr() */

```

```
141  /*
142  * Read bytes per protocol. Filter out zero-length reads.
143  * Die if we don't all the bytes we are supposed to get.
144  */
145  void
146  pread_or_die(int fd,
147               char *buf,
148               int nbytes,
149               void (* diefunc)(int))
150  {
151      if (nbytes > 0)
152      {
153          if (looprw(fd, buf, nbytes, read_no_eintr) != nbytes)
154          {
155              sprintf(
156                  oops_string, "Unable to read %d bytes from fd %d%s",
157                  nbytes, fd, (0 == errno) ? " ", end of file" : "");
158                  oops_msg = oops_string;
159                  (* diefunc)(33);
160              }
161          }
162      } /* end of pread_or_die() */
```

```
165  /*
166  * Read bytes per protocol. Filter out zero-length reads.
167  * Die if we don't all the bytes we are supposed to get.
168  */
169  int
170  pread_or_warn(int fd,
171                char *buf,
172                int nbytes,
173                int (* comm_warning_func)(int))
174  {
175      if (nbytes > 0)
176      {
177          if (looprw(fd, buf, nbytes, read_no_eintr) != nbytes)
178          {
179              sprintf(
180                  oops_string, "Unable to read %d bytes from fd %d%s",
181                  nbytes, fd, (0 == errno) ? " ", end of file" : "");
182                  oops_msg = oops_string;
183                  (* comm_warning_func)(33);
184                  return -1;
185              }
186              return nbytes;
187          }
188      } /* end of pread_or_warn() */
```

```
194  /*
195  * same, but write
196  */
197
198  void
199  pwrite_or_die(int fd,
200               char *buf,
201               int nbytes,
202               void (* diefunc)(int))
203  {
204      if (nbytes > 0)
205      {
206          if (looprw(fd, buf, nbytes, write_no_eintr) != nbytes)
207          {
208              sprintf(oops_string, "Unable to write %d bytes to fd %d",
209                      nbytes, fd);
210              oops_msg = oops_string;
211              (* diefunc)(34);
212          }
213      }
214      /* end of pwrite_or_die() */
```

```
217  /*
218  * same, but write
219  */
220
221  int
222  pwrite_or_warn(int fd,
223                char *buf,
224                int nbytes,
225                int (* comm_warning_func)(int))
226  {
227      if (nbytes > 0)
228      {
229          if (looprw(fd, buf, nbytes, write_no_eintr) != nbytes)
230          {
231              sprintf(oops_string, "Unable to write %d bytes to fd %d",
232                      nbytes, fd);
233              oops_msg = oops_string;
234              (* comm_warning_func)(34);
235              return -1;
236          }
237          return nbytes;
238      }
239      /* end of pwrite_or_warn() */
```

```

241  /*
242  * Send a command packet down an auxproc file descriptor.
243  * Returns number of bytes sent or -1 for unrecoverable error.
244  */
245  int
246  auxcmdpacket(int fd,
247              char cmd,
248              int datalen,
249              char *data)
250  {
251      int retStatus1 = 0;
252      int retStatus2 = 0;
253      int retStatus3 = 0;
254      retStatus1 = write_or_warn(fd, &cmd, 1, auxproc_comm_warning);
255      retStatus2 = write_or_warn(fd, (char *)&datalen, sizeof datalen,
256                                auxproc_comm_warning);
257      retStatus3 = write_or_warn(
258          fd, data, datalen, auxproc_comm_warning);
259      if ((-1 == retStatus1) || (-1 == retStatus2) || (-1 == retStatus3))
260      {
261          return -1;
262      }
263      return (retStatus1 + retStatus2 + retStatus3);
264  }
265  /* end of auxcmdpacket() */

```

```

267  /*
268  * Common portion of auxresults/auxresults_intr
269  */
270  int
271  auxres2(int fd,
272          char cmd,
273          int fixedbuflen,
274          char **fixedbufp)
275  {
276      int datalen;
277      int retStatus1 = 0;
278      int retStatus2 = 0;
279
280      retStatus1 = pread_or_warn(fd, (char *)&datalen, sizeof datalen,
281                                auxproc_comm_warning);
282
283      if (datalen > fixedbuflen)
284      {
285          if ((*fixedbufp = (char *) malloc(datalen)) == NULL)
286          {
287              rbe_log_stats(0, "Could not allocate memory in auxres2");
288              return (-1);
289          }
290      }
291      retStatus2 = pread_or_warn(fd, *fixedbufp, datalen,
292                                auxproc_comm_warning);
293
294      if ((-1 == retStatus1) || (-1 == retStatus2))
295      {
296          return -1;
297      }
298      return datalen;
299  }
300  /* end of auxres2() */
301
302

```



```
305  /*
306  * Read a reply packet beginning with 'cmd';
307  * return results via the buffer passed in fixedbufp,
308  * if it is large enough, else allocate a new buffer.
309  * Returns length of returned results. Fails with -1;
310  */
311
312 int
313 auxresults(int fd,
314            char cmd,
315            int fixedbuflen,
316            char **fixedbufp)
317 {
318     char c;
319     int retStatus1 = 0;
320     int retStatus2 = 0;
321
322     retStatus1 = pread_or_warn(fd, &c, 1, auxproc_comm_warning);
323
324     if (c != cmd)
325     {
326         sprintf(
327             oops_string, "auxresults: expected command '%c',
328                             received command '%c'",
329                             cmd, c);
330         oops_msg = oops_string;
331         errno = 0;
332         auxproc_comm_warning(65);
333
334         /* flush rest of message, if any */
335         retStatus2 = auxres2(fd, cmd, fixedbuflen, fixedbufp);
336
337         return -1; /* report failure */
338     }
339     if (-1 == retStatus1)
340     {
341         return -1;
342     }
343     retStatus2 = auxres2(fd, cmd, fixedbuflen, fixedbufp);
344     return retStatus2;
345     /* end of auxresults() */
}
```

```
350  /* This should not remain an exit, or should it */
351  /* fatal_auxproc should be; indicate unrecoverable
352  auxproc communication error */
353
354  /* Remove use of stderr */
355  /* Pass oops string as input. */
356
357  int
358  auxproc_comm_warning(int tracer_nbr)
359  {
360      char *errno_msg = NULL;
361
362      if (0 != errno)
363      {
364          errno_msg = esl_strerror(errno);
365      }
366      rbe_log_stats(0,
367
368          "%s: UNRECOVERABLE auxproc communication, (
369              progname, tracer_nbr, tracer #&d), %s, %s (%&d)\n",
370              (NULL != oops_msg) ? oops_msg : "",
371              (NULL != errno_msg) ? errno_msg : "", errno);
372      return(tracer_nbr);
373      /* end of auxproc_comm_warning() */
}
```

```

375  /*
376  * Wait, interruptibly,
377  *   for at least one byte to become available on fd.
378  * Returns 1 if at least one byte is available.
379  * Returns 0 if no bytes are available.
380  * Returns -1 for any type of failure, including wait interruption.
381  * Sets errno appropriate when -1 is returned.
382  */
384  int
385  fd_avail_test(int fd)
386  {
387      esl_fset_t rdbits;
388      struct esl_timeval timeout = { 0, 0 };
389      int ret_status;

391      E_FD_ZERO(&rdbits);
392      E_FD_SET(fd, &rdbits);

394      /*
395      * Don't need to examine rdbits after select, since only one
396      * fd is in the set -- therefore return value can be computed
397      * directly from select return value.
398      */
400      ret_status = esl_select(
                    E_FD_SETSIZE, &rdbits, NULL, NULL, &timeout);
402      return ret_status;
404      } /* end of fd_avail_test() */

```

```

409  /*
410  * ChildDone()
411  * Description:
412  * Check for the child being done. If the EINTR error
413  * is encountered for the waitpid then it will be
414  * retried.
415  * Remember that the child will become a defunct process
416  * if its parent process does not wait on it.
417  *
418  * Errno:
419  * NOT ECHILD pid not running or not in group. This is handled
420  * by reset errno to 0 and returning 5.
421  *
422  * EINVAL Invalid argument.
423  * ERANGE The return of the child from waitpid was not expected.
424  *
425  * Parameters:
426  * (I) child_pid -- child pid to check.
427  * (O) child_result -- child's exit status or signal.
428  *   (determined by return value)
429  *
430  * Returns: int
431  * -1 internal error or system error. errno is set.
432  * 0 child still running.
433  * 1 child exited. child_result is exit code.
434  * 2 child signalled (no core), child_result is signal.
435  * 3 child signalled core file generated,
436  *   child_result is signal. *NOT Supported *
437  * 4 child stopped. child_result is signal.
438  * 5 child does not exist or is not a child of the calling process.
439  *
440  * Side effects:
441  * errno can be reset waitpid OR ChildDone.
442  */
443  /* HACK */
444
446  int ChildDone(int child_pid, int *child_result)
447  {
448      int child_exit_val = 0;
449      int wait_status = 0;
451      if (NULL == child_result)
452      {
453          errno = EINVAL;
454          rbe_log_stats(RBRECOVER_MKERR(errno),
455                      "Can't waitpid for child,
456                      ChildDone:invalid parameters.\n");
457      }
458      return -1;
459  }
461  *child_result = 0;
463  #if 0
464      while((0 < (wait_status = waitpid(
465          child_pid, &child_exit_val, WNOHANG))) &&
466          (eintr_retries > 0))
467      {
468          if (EINTR == errno)

```

Page 139 of 144	ChildDone	Thu Jan 03 12:25:21 2008	Page 140 of 144	ChildDone	Thu Jan 03 12:25:21 2008
<pre> 468 3 { 469 3 eintr_retries--; 470 2 } 471 2 else if(ECHILD == errno) 472 3 { 473 3 *child_result = 0; 474 3 errno = 0; 475 3 return 5; 476 2 } 477 2 else 478 3 { 479 3 /* 480 3 * Only retry on EINTR 481 3 */ 482 3 break; 483 2 } 484 1 } 485 1 #endif 487 1 /* 488 1 * Do a waitpid NOHANG. 489 1 * Set ECHILD to return 5. 490 1 * longer around. 491 1 */ 493 1 do 494 2 { 495 2 wait_status = waitpid(child_pid, &child_exit_val, WNOHANG); 497 1 } while((-1 == wait_status) && (EINTR == errno)); 500 1 switch(wait_status) 501 2 { 502 2 case(-1): /* error encountered */ 504 2 if(ECHILD == errno) 505 3 { 506 3 /* Lets handle the no child case as not an internal error. 507 3 *child_result = 0; 508 3 errno = 0; 509 3 return 5; 510 2 */ 511 2 rbe_log_stats(RBRECOVER_MKERR(errno), 512 2 "Can't waitpid for child, error %s == %d\n", 513 2 strerror(errno), errno); 515 2 return -1; 517 2 case(0): /* child still running */ 519 2 return 0; 521 2 default: /* child is NOT running */ 523 2 if (WIFEXITED(child_exit_val)) 524 3 { 525 3 *child_result = WEXITSTATUS(child_exit_val); 526 3 return 1; 527 2 } 528 2 else if (WIFSIGNALED(child_exit_val)) 529 3 { 530 3 *child_result = WTERMSIG(child_exit_val); </pre>			<pre> 531 3 #if 0 532 3 /* WCOREDUMP is not _POSIX_C_SOURCE */ 533 3 if(WCOREDUMP(child_exit_val)) 534 3 return 3; 535 3 else 536 3 return 2; 537 3 #else 538 3 return 2; 539 3 #endif 540 2 } 541 2 else if (WIFSTOPPED(child_exit_val)) 542 3 { 543 3 *child_result = WSTOPSIG(child_exit_val); 544 3 return 4; 545 2 } 546 2 else 547 3 { 548 3 errno = ERANGE; 549 3 rbe_log_stats(RBRECOVER_MKERR(errno), 550 3 "Can't determine waitpid status of the 551 3 child process.\n"); 552 2 return -1; 553 1 } 554 1 } /* End ChildDone() */ </pre>		
Page 139 of 144	RSlauxSupp.c 15	Thu Jan 03 12:25:21 2008	Page 140 of 144	RSlauxSupp.c 16	Thu Jan 03 12:25:21 2008

